ENTATION PAGE

PORT DATE	3. REPORT TYPE AND DATES COVERED
	FINAL 1 Nov 89 TO 30 Apr 92

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
RECEPTORAL AND NEURAL ALIASING	AFOSR-89-0126
	61102F
6. AUTHOR(S)	2313
Dr Robert Allen Smith	A5

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
University of New Hampshire 99 Madbry Road Durham, NH 03824	AFOSR-TR· ·· 3 0 0 6 0

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING / MONITORING AGENCY REPORT NUMBER
Dr John F. Tangney AFOSR/NL 110 Duncan Avenue, Suite B115 Bolling AFB DC 20332-0001	

11. SUPPLEMENTARY NOTES

DTIC
ELECTE
MAR 0 2 1993

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

13. ABSTRACT (Maximum 200 words)

This project saw some unexpected successes, and some even more unexpected failures. Chief among the latter was our total failure to replicate Williams' observations of receptoral aliasing, either in the fovea or the parafovea. Despite much communication between us and Williams, no explanation of this discrepancy has been found. We present substantial evidence below to eliminate the possibility of technical failure in our experiments. Unable to study aliasing directly, we pursued a very different approach to the problems of retinal geometry and aliasing: we studied methods to directly map the parafoveal visual field into its constituent summation areas (presumably receptive fields). Although we have not fully solved the very difficult problem of eye-movements, several successful studies were performed, notably; (1) the development of a precise new method for measuring fixation accuracy using afterimages, and (2) studies of spatial summation for isoluminant spots in the parafovea. In addition, a computer model was designed that produces sample ganglion cell lattices and models the process of chromatic identification. Finally, we diverged considerably from the original focus of the project to design what we believe to be the most powerful stimulus display for visual research currently available. This is now a commercial product, showing both scientific and financial success.

14. SUBJECT TERMS		15. NUMBER OF PAGES
		16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT (U)	18. SECURITY CLASSIFICATION OF THIS PAGE (U)	19. SECURITY CLASSIFICATION OF ABSTRACT (U)	20. LIMITATION OF ABSTRACT (U)

HANS-LUKAS TEUBER

VISION LABORATORY

99 Madbury Road
Durham, NH 03824

RECEPTORAL AND NEURAL ALIASING

Robert Allen Smith, PhD
Principal Investigator

January 30, 1993

Final Report: Period 11/1/89 - 4/30/92

Grant: AFOSR-90-0126 A

DISCLAIMER NOTICE

HANS-LUKAS TEUBER

VISION LABORATORY

99 Madbury Road
Durham, NH 03824

RECEPTORAL AND NEURAL ALIASING

Robert Allen Smith, PhD
Principal Investigator

January 30, 1993

Final Report: Period 11/1/89 - 4/30/92

Grant: AFOSR-90-0126 A

TABLE OF CONTENTS

SUMMARY

This project saw some unexpected successes, and some even more unexpected failures. Chief among the latter was our total failure to replicate Williams' observations of receptoral aliasing, either in the fovea or the parafovea. Despite much communication between us and Williams, no explanation of this discrepancy has been found. We present substantial evidence below to eliminate the possibility of technical failure in our experiments. Unable to study aliasing directly, we pursued a very different approach to the problems of retinal geometry and aliasing: we studied methods to directly map the parafoveal visual field into its constituent summation areas (presumably receptive fields). Although we have not fully solved the very difficult problem of eye-movements, several successful studies were performed, notably; (1) the development of a precise new method for measuring fixation accuracy using afterimages, and (2) studies of spatial summation for isoluminant spots in the parafovea. In addition, a computer model was designed that produces sample ganglion cell lattices and models the process of chromatic identification. Finally, we diverged considerably from the original focus of the project to design what we believe to be the most powerful stimulus display for visual research currently available. This is now a commercial product, showing both scientific and financial success.

2

I. *OUTLINE OF RESEARCH ACTIVITY*

A. Receptoral Aliasing and the Optical Bench

The first objective was the construction of the laser interferometer. This was begun in January '89, and proceeded rapidly. It was clear within several weeks that we could project fringes on the retina, and that problems of vibration and abberation were not major issues. The design of the interferometer is a new one, provided to us by Dave Williams, who was constructing a similar device at the same time. This design offers a distinct advantage over previous interferometers in that the two interfering beams traverse the same optical path, only in opposite directions. This cancels much of the effect of vibration.

Several problems surround the interferometer. The interferometer depends crucially on the ability of its polarizing beam-splitter cubes to separate the beam into opposing polarizations with >99.9% purity. Both we and Dave Williams became aware simultaneously that the cubes used fall far short of their specifications in this regard. The result of this crosstalk between the polarized beams is that we cannot reduce the fringe contrast below about 4%. Although we devised several possible solutions, the issue did not present an actual limitation on our research and so was of no practical consequence.

A much more serious problem is apparently (after much study) not of a technical nature. A central assumption of our proposal was that, having duplicated Williams' apparatus, we would readily replicate, and then extend, his reported findings. This never happened. The original discrepancy between our respective results (predating this proposal) was that we found that orientation discrimination at 7 degrees parafoveal ceased at about 10 c/deg, while Williams reported discrimination to about 25 c/deg. The issue is an important one since the discrimination limit is the operational definition of the onset of aliasing. Thus our subjects appear to be observing with a lattice with about twice the spacing of Williams'. There is still no good explanation for this discrepancy.

To our considerable surprise, the use of interference fringes not only failed to eliminate the above discrepancy, but we found that we could not even replicate Williams' published results in foveal viewing ("zebra stripes"). Above about 60 c/deg our fringes become undetectable and remain so at all higher frequencies. We are not alone in this; at least one other group has seriously tried to replicate Williams and failed (personal communication: researcher wishes to remain anonymous). We have discussed this many times and at length with Williams. The only explanation which was forthcoming was the possibility of unobserved high-frequency, low-amplitude vibrations in our system. The problem with such vibrations is that they could be quite unobservable at low frequencies, yet a vibration of the pattern whose amplitude is comparable to the actual fringe spacing (NOT the observed spacing, which may be 2 orders of magnitude larger) would reduce the fringe contrast (and that of the observed alias) to near zero. Having recently acquired a graduate student who is quite unable to see aliasing, Williams is now much more receptive to the idea of individual differences. The nature of these, however, remains entirely obscure.

Possible sources of such vibration fall into two classes; 1) intrinsic to

3

the apparatus, and 2) failure of vibration isolation. Intrinsic vibrations would emanate from electrical devices on the optical table. These include 1) the laser, 2) the acousto-optic modulators, and 3) the stepper motors which adjust spatial frequency and orientation. The first two of these cannot readily be changed, since they are essential to the operation of the apparatus. Lasers, however, are not considered to be a source of vibration, and the a/o modulators produce infinitesimal vibrations at 40 Mhz; such vibrations would not propagate far from their source. The stepper motors, however, are energized continuously at 60 Hz, and are directly connected to beam-deflecting mirrors; they seem a prime candidate. However, this vibration is readily eliminated by unplugging the steppers and making adjustments manually. This had no effect on our results.

The obvious path to vibration isolation -- purchase of an isolated table -- was not economically feasible. Williams solved this problem by using a very rigid support and suspending the actual table on four motor-scooter inner-tubes. Our original apparatus used a mattress to isolate the table; perhaps this is inadequate. We have therefore gone to a two-stage vibration filter. The first stage is the mattress with a 400 pound steel plate resting on it. This stage has a resonant frequency of about 2 Hz, which is unusually good for a small table (24" x 48"). The second stage is much lighter (decoupling it from the first). It consists of four inner-tubes resting on the steel plate and supporting a 2" plywood sheet which holds the actual optical table. Finally the whole apparatus rests on a slab floor. Thus our isolation system would seem to be at least as effective as Williams'.

Despite all of the above efforts, our results remained unchanged. At this point we doubted whether vibration was in fact the cause of our difficulties. We devised the following test of the apparatus, which shows definitely that high-frequency fringes are being produced. There is no reason why the fringes must be projected on the retina. Instead we projected them directly on the sensitive element of a CCD camera. By moving the camera to larger distances the fringes could be made any convenient size. Figure 1 shows fringes which would have had spatial frequencies of 10 c/deg and 150 c/deg on the retina. The 10 c/deg grating is clearly visible; the image of the 150 c/deg grating, however, is an alias! The Nyquist limit for our CCD camera is about 80 c/deg. Note that the contrast of the alias is little reduced over that of the low frequency fringe; this can happen *only* if the contrast of the 150 c/deg fringe remains high. Clearly, then, high frequency gratings are not being lost to vibration or any other cause.

Although it has not reproduced clearly in Figure 1, there is a lot of diffraction noise in the image, presumably from scattering by imperfections in the optics. We have done many thorough cleanings of the optics, to try to reduce this. In fact, however, this noise problem is quite small compared to wht happens when the gratings are viewed by the living eye. In this case, every piece of detritus in the aqueous humor, and every irregularity on any of the optical surfaces of the eye produces a clear image on the retina. This is normal when viewing by coherent light. The noise level from these sources appears substantial, but there is no obvious way to actually measure it since it is entirely entoptic; perhaps the more subtle Moire' fringes are being totally masked. Of the three noisy optical surfaces in the eye (the corneal surface, and the two lens surfaces) it is possible to eliminate one by focussing the laser

4

Figure 1   Interference fringes projected onto a CCD camera. Left side: 10 cpd spatial frequency at plane of retina. Right side: alias of 150 cpd.

beam directly on that surface. The corneal surface is the dirtiest, so we focus on this one, but the improvement is by no means dramatic.

At this point, our efforts in this direction have met with a dead stymie. If we are unable to observe aliasing directly (i.e. "zebra stripes") then there is little we can do which we have not done previously with ordinary displays. We reluctantly left the study of aliasing with interference fringes for other related investigations.

B. Mapping the Retinal Mosaic: the problem of eye movements

This overall project deals with a variety of indirect methods for studying the retinal mosaic, either receptoral or neural. All of our previous work, however, approaches the problem from the spatial frequency aspect (e.g. the measurement of Nyquist frequencies.) Given the difficulties described above, we decided it would be useful to work in the spatial domain, by directly mapping the functional, retinal mosaic using the detection of small points of light. The foveal cone lattice is too dense to map relative to the large point-spread function (Jennings and Charman, 1981), but in the parafovea the lattice density falls very sharply, while optical quality is only slightly reduced. We recently (Smith and Cass, 19899) published data on spatial summation in the parafovea, which suggests that summation aea are about 3' in diameter. Our measurements of the Nyquist limit (Smith and Cass, 1987) offer a direct measurement of parafoveal lattice spacing which is also about 3'. While these figures are certainly imprecise, they are accurate enough to suggest two things. 1) Both the summation area and lattice spacing are considerably larger that the point-spread function (about 1.2'), suggesting it might be possible to probe these structures with tiny points of light. 2) Since their separation and diameters are about the same, summation areas will be relatively non-overlapping; thus their individual properties might be detectable by such probes. On the other hand, the parafoveal cone mosaic -- with cone diameters of 1.5' and separations of 2.3'(Hirsch and

5

Miller, 1987) -- is probably too fine for probing in this way.

These conclusions are summarized in the following table.

TABLE 1

Entity	Eccentricity	Diameter	Separation
Sum. Area	7'	2.9'	4.2'
Cone	1'	0.8'	1.1'
	6'	1.5'	2.28'

Initial tests on the detection of small points of white light were not encouraging, so we devised a more sensitive test based on the observation that the majority of ganglion cells (the midget ganglion cells, about 85%) are opponent-color-tuned. Changing the task to the discrimination of the color of a spot (which was one of two opponent colors), might effectively "thin out" the mosaic, since such a spot would activate only areas sensitive to its particular colour

Our initial experiments were done with very careful fixation, and voluminous data collection to average out positional uncertainty. The major substantive conclusion from these studies was that there was a numerically large (and highly significant) negative correlation between areas of red and of green sensitivity. Blue vs. yellow shows an equally large effect. This clearly shows that areas sensitive to opponent colors occupy largely separate and non-overlapping retinal areas. (Note that there is no significant correlation between areas sensitive to red and blue, or to green and yellow.) On the other hand, the methodological conclusion from these studies is that the noise resulting from uncontrolled eye movements precludes making the kind of detailed sensitivity map we hope to achieve; a 6x6 array of detection probabilities (the largest we found practical) hardly constitutes a map. We therefore embarked on the design of an eye-position monitor. In the '89 report we envisioned a video system, which would monitor both the pupil and the first Purkinje reflection. Both of these will shift equally with head position, but the Purkinje reflection shifts somewhat less with rotation of the eye than does the pupil. Thus a differential measurement should correlate with eye position.

1. Fixation Error

The major difficulty in mapping such small areas of the retina is, of course, the inability of a subject to maintain perfect fixation. We set up conditions to provide maximal assistance to the subject: the subject had a good fixation target (a cross within a square), and was allowed to trigger trials only when he believed his fixation to be very good. There do not seem to be data which clearly test the adequacy of this technique. Estimates of fixation accuracy over relatively short intervals (e.g. Eizenman et al, 1985) suggest that actual

6

fixation position varies between +/- 10 arc-minutes, with a standard deviation of about 5 - 7 arc-minutes. In our paradigm, fixation should be at least this good, but it might be considerably better. Furthermore, any study relying on instruments to make fine measurements of eye position are inevitably subject to questions of calibration, etc. We devised a purely psychophysical technique for measuring fixation error in conditions similar to our experiment. The technique relies on after-images and human observers' hyperacuity at making positional judgments.

An after-image of a short vertical line was produced by looking into a strobe-light covered with a mask. A fixation point appeared on a crt screen, followed by two vertically oriented flanking lines, twenty pixels apart, for a duration of 200 msec. The observer's job was to rate the position of the after-image relative to the flanking lines. Control data were also taken, substituting an actual line in place of the after-image.

The results for the control data indicated that untrained subjects can perform the task with a standard error of about .7 arc-minutes. Since fixation error is presumably several arc-minutes, the positional judgments have more than enough accuracy to measure fixation error.

Three experimental sessions were run with observer DJS. The standard deviation of the after-image placement relative to the flanking lines ranged between 3.8 and 5.1 minutes. While these data indicate that fixation accuracy is good enough for many purposes, it may not be accurate enough for mapping small areas of the parafovea where receptive field sizes of ganglion cells are as small as 3 arc-minutes (Derrington & Lennie, 1984). We decided that precise mapping required greater fixation accuracy, and that the best way to achieve this was with a high-precision eye-monitor.

2. Eye-Position Monitor

A difficulty in this work is that we have only the crudest means of estimating the true direction of gaze, to compare with the predictions of our eye-position monitor. Initially we simply ran our colour-discrimination experiments using the eye position monitor; any improvement in the data would have reflected a reduction in eye-position inaccuracy due to the monitor. There was no improvement. We therefore turned to more basic studies of the monitor's behavior. First we considered the accuracy of the routines which localize the pupil and the reflection. For our desired accuracy (1'-2') it will be necessary to localize to about 4 microns or 0.03x the size of a pixel. Our localization algorithms are designed to perform in the range of hyperacuity, and various experiments indicate that the routines are indeed accurate enough.

a. Artificial Eye Data

The most direct experiment was done with an artificial eye moved by a micrometer. An index card was produced with a black circle to represent the pupil. A white circle within the black circle represented the reflection of the

light source off the surface of the cornea (the *corneal reflux*). The micrometer
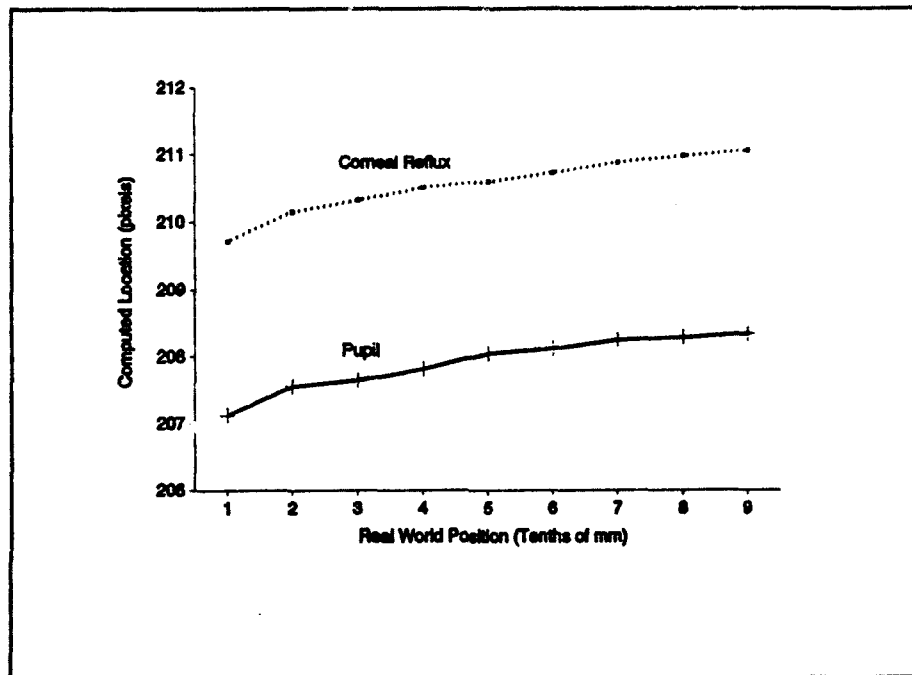was used to move the artificial eye in 0.01 mm steps. The raw data are presented
in Figure 2.

Figure 2   Results of experiment with artificial eye.

The most valid analysis of the results is a comparison of the difference
between the computed positions of the artificial pupil and reflux, since this
will be our actual experimental measure.   The standard deviation of the
difference was computed to be 4 arc-minutes, which we take as a measure of the
resolution of the eye-monitor under ideal circumstances.   Ultimately we would
like to refine our algorithms to improve this accuracy by about 2x, but presently
there are far more important limitations to be dealt with.   Note that this
measurement equates to .058 pixels--demonstrating that the monitor is operating
.11 into the *hyper-acuity* range.

b.   Natural eye data

We then proceeded to make measurements of large excursions of a real eye,
where we could trust the observer's report of fixation.   We found that the
differential movement of the corneal reflection and the pupil was simply not
great enough to provide the needed accuracy.   Probably the situation would be
improved by using higher Purkinje images, as in the Stanford Eye Tracker, but

8

these are faint and usually lost within the large first image. Surprisingly, if we greatly reduce the size of the first image, making the higher images more visible, then our hyper-acuity algorithms lose their power and accuracy declines. We decided instead to try to control for head motion directly, by placing a small, accurately resolvable, target in the canthus of the eye. Small styrofoam beads worked well for this. Figure 3 shows a typical example. (Note that the image processor splits the image, leaving out the unimportant part of the eye between the pupil and the canthus.)

Figure 3   A captured video frame showing the pupil and corneal reflux (on right) and the styrofoam ball (on left).

It is clear that a complex mathematical relation must relate the observed quantities (primarily the centers of the pupil and the bead) with direction of gaze. To determine this relation *a priori* is probably not possible. In general, the camera cannot observe the eye exactly along the line of sight, so the image coordinates are not orthogonal in the plane of the eye, nor are they related one-to-one to the cardinal directions of gaze. Similarly there is no practical way to calculate the relation between pupil center and direction of gaze (even with head position controlled). Because the angles involved are all very small, however, we may assume that the prediction equation is linear. This allows us to use a form of learning paradigm, in which the computer constructs the prediction equation by linear regression on a set of known eye-movements. In practice, we supplied the subject with four fixation points at the corners of a square; initially we used a 10' square, but found that we required a one degree square. The subject's eye position was recorded several times while fixating each point, to produce a regression data set.

The linear regression was much more than a way to determine an equation, however. By detailed analysis of the individual correlations, we gained

9

Figure 4 The predicted values are the result of a regression analysis based on the tracker's computed values of pupil and reflux location. The units are display-pixels; about 0.6' of visual angle.

considerable insight into our sources of error. For example, the observation that trial number had a high correlation with other observables helped us identify a drift problem. At this point, the major obstacle is with the styrofoam bead as an index of head position. Under normal conditions, the location of the bead has a large standard error, and very little correlation with anything. It is readily observed that the bead moves with eye blinks; possibly it is not returning accurately to its initial position. This is difficult to verify, but an experiment in which the subject was extremely careful not to blink or move his head yielded noticeably better predictions, but the best standard errors of the estimate are still on the order of 20 - 30 arc- minutes (Figure 4). A portion of the standard error is due to genuine fixation error, but as shown in Section B.1, fixation error is probably well under 10 arc-minutes, and therefore not the major component of the tracker's inaccuracy. The position of the bead, however, made almost no contribution to the prediction. This all implies that bead position is too unreliable to use as a control for head movement. Our current thinking is that a bead fastened to the bridge of the nose may be more stable. Unfortunately this will require a second video camera: before proceeding to this expense we are seeking ways to demonstrate the viability of this approach

10

with our present apparatus.

C. Chromatic Summation Areas

The entire enterprise of mapping the ganglion cell mosaic with psychophysical techniques is based on certain assumptions that are derived primarily from physiological studies performed on macaques. We decided that psychophysical data would be useful to supplement the physiological data, particularly as it related to our assumption of receptive field sizes. We therefore embarked on a study of chromatic summation areas in the periphery (Smith & Swift, 1991).

All experiments were performed on a CRT. Red spots of light were flashed for 200 msec on an equiluminant green background at varying eccentricities. The other major independent variable was spot size. The dependent variable in all cases was chromatic contrast. Functional equiluminance was determined by measuring detection thresholds at various red-green ratios, and picking the ratio that yielded the highest threshold.


EQUILUMINANT SPOTS

Figure 5   Results of Chromatic Summation Experiment. Note that Riccc's Law implies a zero slope; full summation implies slope = 2; data fall in-between (Piper's Law)

If a stimulus is within the receptive field of a detector, then Ricco's Law should obtain; i.e. there should be a complete trade-off between area and contrast. Another way of phrasing this is that total energy at threshold should not change as a function of spot size--yielding flat curves. No summation, on the other hand, requires a slope of two, since energy rises as the square of spot diameter. The data (Figure 5) are clearly in between the two extremes. This relationship has been referred to as *Piper's Law*, although the underlying mechanism is not well understood. Presumably then, the chromatic summation areas are less than 14 arc-minutes--the smallest spot sizes in the experiment. The limiting factor on spot size was not hardware resolution; rather it was detection threshold, which was close to 100% contrast for spot sizes of 10 arc-minutes. These results raise the interesting question of whether the response of a single ganglion cell to a near optimal stimulus is sufficient for detection.

A second experiment was done with luminance spots in order to compare the results to those with chromatic stimuli (Figure 6). Note that while thresholds are lower, once again Ricco's Law did not obtain.

LUMINANCE SPOTS

Figure 6   Results of summation experiment with luminance contrast.

D. A Model for Chromatic Identification

In the course of trying to map the retinal mosaic it occurred to us that we were dealing with many unspoken assumptions about underlying processes. It struck us that it would be better science to start with a working model of chromatic identification that would explicitly state the assumptions we were working under, and also allow us to test the implications of the various assumptions. In other words, we could freely vary parameters in the model to test the robustness of the assumptions.

1. Available Physiological Data

From such sources as Zrenner *et al*, 1989, Derrington & Lennie, 1984, Livingston & Hubel, 1987, and Perry & Cowey, 1985, we have made the following assumptions regarding the ganglion cells at 7° in the periphery.

- about 85% of the ganglion cells project to parvocellular layers in the lateral geniculate; the other 15 % project to the magnocellular layers

- over 80% of the parvocellular cells are chromatically opponent; the magnocellular cells are not considered to be color-opponent

- there is considerable variability in the neutral point for the 80% of the parvo cells that are color-opponent

- the majority of the color opponent cells are red-green rather than blue-yellow

- receptive field centers of parvo cells at 7° in the periphery are about 5 arc-minutes with relatively little variability

- spacing between ganglion cell centers is also approximately 5 arc-minutes at 7° in the periphery

2. Assumptions Based on Known Physiology

The facts stated above, along with other commonly held beliefs lead to further assumptions related to chromatic identification. Some of the assumptions are fairly tenuous, but this is not a weakness of the model. Rather the weaker assumptions can be identified, modified and examined in the model.

- *Red on-center and green on-center cells have non-overlapping receptive fields.* Since there is close to a one-to-one correspondence between cones and ganglion cells and since receptive field sizes of ganglion cells and cones at 7° in the periphery match reasonably well, it is likely that on-centers receive inputs from only one or two cones. Since cones don't overlap, neither should ganglion cell receptive fields

13

- *The distribution of ganglion cell types is completely random.* This assumption is not so much based on known physiology as it is on lack of evidence to the contrary. It would require some active mechanism to achieve some regularity in the distribution of ganglion cells, and no such mechanism has ever been reported. This turns out to be a fairly important assumption in the model, and one that will have to be carefully evaluated.

- *The likelihood of a red response to a spot of light is determined by the responses of the red on-center ganglion cells minus the responses of the green on-center ganglion cells.* This is seldom stated explicitly but it follows from many commonly held beliefs. A more thorough analysis of signal and noise might suggest a more subtle criterion.

3. A Description of the Model

There are two stages in the model: (1) construction of a representative ganglion cell mosaic, and (2) simulation of an experiment that presents small spots of colored lights.

a. Constructing a Representative Mosaic

A Poisson process is first used to place as many ganglion cells as can comfortably fit into a hexagonal area without overlap. All additional cells are then added without the requirement for non-overlap. Cells are randomly placed according to the following proportions: 35% red on-center, 35% green, 17.5% blue, 17.5% yellow. Non-opponent parvo cells and magno cells are ignored since they presumably do not contribute to chromatic identification. The computer then performs an iterative *pushing* algorithm that moves the additional cells around to achieve a reasonable degree of non-overlap. Although the original area is hexagonal, the final packing turns out to be only crudely hexagonal, which is probably appropriate for ganglion cells. A sample mosaic is presented in Figure 7. Note the clustering of red and green on-center cells. This, while slightly counter-intuitive, is solely a product of random chance.

b. Simulation Results

The model simulates the following experiment: Sixteen different fixation positions are use, randomly, over the course of the experiment. At each fixation position, there are 15 random presentations of a red spot of light, and 15 of a green spot. The observer's job is to respond *red* or *green*. In a second condition the observer may respond *red*, *green* or *neutral*. The simulation includes a value for fixation error (See page 6), so that the actual spot of light is placed at some random gaussian distance from the fixation point. The response of the ganglion cells is determined by the sensitivity of a particular type of cell to the color of the stimulus spot and by the gaussian distance of the spot from the receptive field center. The gaussian distance factor controls both for the gaussian nature of receptive field centers and the gaussian nature
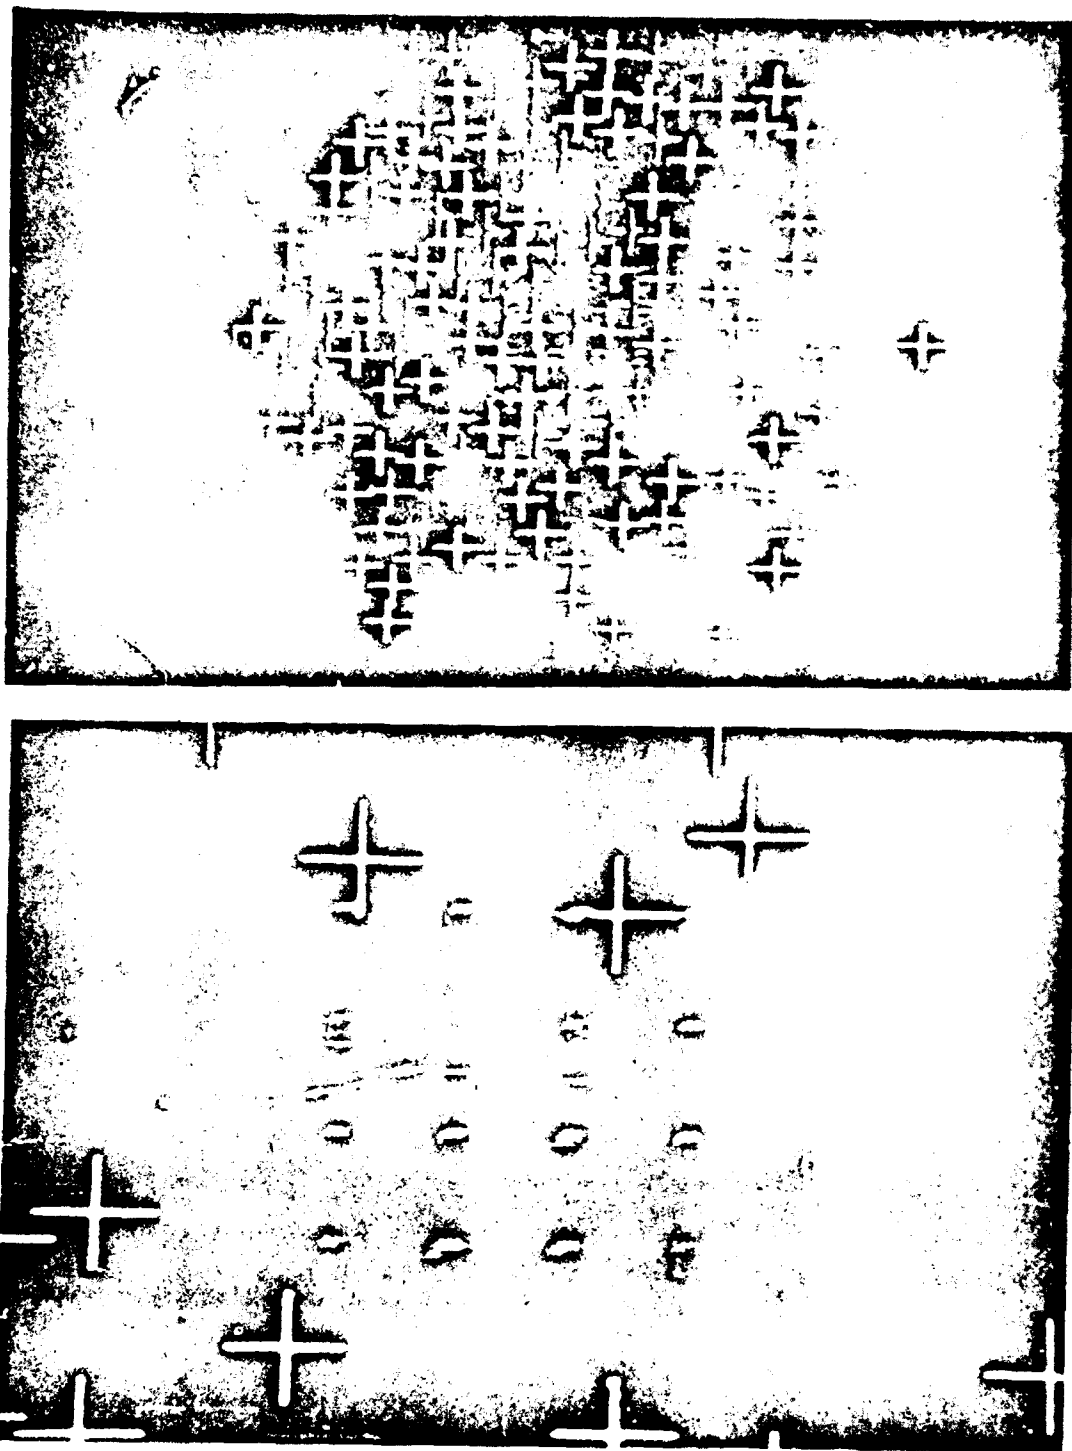
14

Figure 7 (Top) A sample array of color-opponent ganglion cells. The white, green and red blotches represent the corresponding incorrect responses at those fixation positions. (Bottom) Close-up of results.

15

of the point-spread function of the optics of the eye. Gaussian noise is then added to the response. Finally, if the net result is greater than 0, the appropriate response is given; otherwise the incorrect response is given. In the three-alternative response condition, a threshold is added to the decision rule. If the absolute value of the result is below threshold, then the response is *neutral*.

Of major concern in the simulation was the question of whether fixation error would be too large to allow for significant results. Fixation error was therefore varie parametrically and Chi-Square analyses were performed to test the following H_1: Are the relative numbers of *red, green,* and *neutral* responses independent of fixation position. The following table shows the results of the analysis for various fixation standard deviations.

$\sigma_{fixation}$ (arc-minutes)	2-alternatives	3-alternatives
	Typical *p* value	
4.2	.001	.001
6.3	.01	.01
8.4	.1	.01
10.4	>.1	>.1
12.5	>.1	>.1
14.6	>.1	>.1

It can be readily seen that statistically significant results should be obtainable with standard deviations for fixation error up to 8 arc-minutes, comfortably within the range reported on page 6. While we have obtained statistically significant results in this experiment (Smith & Swift, 1990), it has not been as easy as the above analysis would suggest. We suspect that the luminance levels of our CRT may not be quite sufficient for the task. We have purchased appropriate equipment for using external light sources under computer control, and will report on the results of these experiments in next year's annual report. We are also in the process of conducting simulations with a variety of assumptions, and hope to achieve a more complete match between the actual and simulated data.

E. The VisionWorks Visual Stimulus Generator

During the tenure of this grant, the PI was also significantly involved in the founding of Vision Research Graphics, Inc., whose purpose was to develop state of the art, computer-graphics-based visual stimulators for use by vision researchers, and ultimately by the ophthalmological community. With the concurrence of the AFOSR program director, a significant part of the PI's time under this grant was devoted to the design of VRG's visual stimulators. Although it is quite impossible to disentangle exactly what portions of this effort

may have been sponsored by AFOSR, the results have been impressive, and certain parts can be clearly pointed out as the result of the PI's efforts. VRG has two basic product-lines, 1) pcStereoscope -- a PC-based stereoscope of modest cost and wide applicability, and 2) VisionWorks -- a major integrated software/hardware package which we explicitly intend to be the finest visual stimulator available.

1. pcStereoscope

The pcStereoscope(tm) is a low-cost computer-graphics stereoscope which operates on an IBM/PC or compatible computer with an EGA or VGA display. It utilizes a pair of LCD shutter-goggles to present alternate frames of the CRT display to opposite eyes. The software maintains two distinct images of the visual scene in two hardware pages of the EGA: one for each eye. This permits stereo viewing (when the pages are a stereo pair) as well as haploscopic viewing of any arbitrary sort, since the images to the two eyes are entirely independent. pcStereocsope is being used as a stereo display in about 2 dozen laboratories across the world, and a variety of studies have actually utilized it. We have presented three papers which used the pcStereoscope. At least one of these -- a comparison of four distinct measures of the ability to fuse stereo images -- could not have been done with any other apparatus, since the pcStereoscope could generate all four paradigms with identical luminances, contrasts, etc, and all in a single sitting. In an experimental setting, the pcStereoscope is supported by its Software Development System. The low-level basics of this graphics library (e.g drawing of points, lines, etc.) were written by the author, and -- at the time -- were determined to be the fastest available for the EGA.

The primary use of the pcStereoscope, however, has been in implementing VisionLab(tm) -- an upper-level college visual perception course built around some three dozen demonstrations, and seven actual experiments. The demonstrations include all of the well-known visual demonstrations and illusions, as well as some unusual ones, and a few which were discovered with the pcStereoscope. (One of the latter is the demonstration of stereopsis from pure form (form-from-motion, form-from-subjective contours, etc) without any true binocular cues. This was described in detail in our first Annual Report.) Our demonstrations are unique among computerized visual demonstration programs, in that the student can vary every reasonable parameter of the display. For example, in the Mueller-Lyer illusion the student can change all the dimensions of the lines, as well as the angle of the "arrowsheads". Using our stereo facilities, the student can place the arrowheads in a different depth plane from the shafts, or he can vary the angle of the arrowheads in depth, so that they resemble 3-d corners. Finally all of the above can be done with the illusion in a random-dot stereogram form, all monocular cues being eliminated.

The experiments, though fewer in number play an important role in that 1) they show that visual science is not just clever demonstrations, and 2) they actually provide a course in psychophysical methodology. We offer seven psychophysical methods, from simple method-of-adjustment to signal detection analysis. Thus in the Mueller-Lyer exepriment (an elaboration of the demonstration described above) the student can observe the differences and trade-offs between simple MOA nulling of the illusion, versus forced-choice, method of

17

constant stimuli, etc. Having selected the most appropriate psychophysical method, he can then measure the effects of the various stimulus manipulations described.

A specialized application of the pcStereoscope is Stereo-Snapshot(tm), which was developed entirely by the author. This is a TSR which remains in the PC while the user runs any program capable of generating a stereo pair of images; we are specifically targetting the 3-d CAD market, since such programs can produce a stereo pair by a simple rotation about the vertical axis. The user generates each member of the stereo pair, and saves it to disk with Stereo-Snapshot. Then a call to Stereo-Snapshot displays the pair in stereo. Note that there is no alteration to the host program, and being a TSR, Stereo-Snapshot operates without ever exiting the host. Thus Stereo-Snapshot works with essentially any 3-d CAD or 3-d studio-type program.

2. VisionWorks(tm)

VisionWorks is VRG's high-level visual stimulus generator. It consists of an integrated package of hardware and software designed so that many experiments can be designed entirely from a menu interface, while unusual experiments can be produced with a minimum of actual programming. At this point, such custom programming is more complex than we would like, and we are working on a more automated, user-friendly system. Nonetheless, the majority of applications which we write for customers are done entirely from menu selections. (The only reason we write such simple applications is that we include a certain amount of custom programming with each system; a new customer generally doesn't realize how simple his proposed experiment is!)

Viewed from low-level to high-level, *VisionWorks* consists of the hardware and operating system, the software development environment, and the *VisionWorks Manager*. The operating system provides the functional link which ties the hardware together. The software development environment is a collection of tools that allows the user to efficiently program the hardware. The *VisionWorks Manager* is code written by *VRG* which streamlines the production and management of vision experiments. Each level is dependent upon the previous one. However, the converse is not true. The software development environment can be used without the *VisionWorks Manager*, and often will be.

VisionWorks is supplied as a base system with optional modules. The base system consists of hardware for a standard (non-stereoscopic) display system, the software development environment (C++ compiler, *TIGA* software driver, *MenuMaker* utility and the *VRG* function library), and the *VisionWorks Manager*. The *VisionWorks Manager* includes facilities for experiment data output and for running preconfigured experiments in a simplified manner. The latter functionality will be referred to as the *Experiment Manager*. The *StimulusMaker* and *MethodMaker* (Psychophysics) Modules are optional.

VisionWorks modules are bodies of code which are written to conventions defined by a VisionWorks standard. Modules perform specific functions, such as stimulus generation, but are written in a manner which allows them to be linked together within the framework of a *VisionWorks* application. Modules written by

18

VRG are large and extremely versatile. Modules can also be written by users, and will be smaller and more specialized.

VisionWorks applications are stand-alone programs, generally experiments, that are produced by users or by *VRG*. Applications will often use modules written by *VRG*, such as the *StimulusMaker* and *MethodMaker* modules. A specific application will typically not make use of all the functionality of a *VRG* written module. Rather, the application will incorporate only the features it needs from a module.

VisionWorks has 3 color modes: normal RGB, gamma RGB, and gamma monochrome. Normal RGB mode is a standard color mode with 256 color levels available for each color gun. Gamma_RGB mode is a gamma-corrected color mode in which a user will specify a color value within a range of 0-255 and the output from the color gun will be the closest value to linear within that range.

Gamma monochrome is a special mode which operates only with a monochrome monitor and a VRG grey-scale expander. The grey-scale expander takes the output of the three color-DACs and adds them with different weights, so that the red DAC supplies the course luminance, while the green and blue DACs are progressively finer adjustments. While such a system has been used before, it is subject to the major problem that DACS are typically not accurate to much better than their smallest step. Thus while it might seem that three 8-bit DACs, suitably weighted, would provide 24 bit resolution, the actual accuracy is only 9 or 10 bits. We have circumvented this problem by actually measuring the output of each DAC step to 15 bits accuracy. A computer program produces from these data a 32k*3byte look-up table which takes a 15 bit brightness value and returns 3 bytes to be sent to the 3 DACs. Monotonicity is guaranteed to 15 bits, though we are able to measure actual luminance linearity to only 13 bits for practical reasons. 13 bits linearity seems accurate enough for any purpose, however. To achieve this, it is necessary to compensate for the considerable luminance non-linearity in the CRT phosphors. This is accomplished with the VRG Photometric Linearization System. This is a high-precision photometer (designed by the author) coupled to a major program which controls the entire calibration procedure. With only minor operator intervention, this system accomplishes all of the calibrations described above in about 1/2 hour; done by hand these would require a full day. All new systems are calibrated before being shipped; in addition the Calibration System is available as an option to experimenters requiring highly reliable calibration.

The two most important parts of the *VisionWorks* software are StimulusMaker and MethodMaker. StimulusMaker is a menu-driven stimulus-producer, which frees the experimenter from having to program any of the commonly-used visual stimuli. A complete description of this module (which represents about a man-year's effort) is beyond the scope of this document; Appendix 1 presents a simple list (running to 7 pages!) of the menu options. MethodMaker is an elaboration of the idea of making the psychophysical methods interchangeable which was introduced in VisionLab. This currently makes about 10 psychophysical methods available to the investigator, including such complex techniques as the QUEST staircase and a full signal-detection analysis. Freedom from having to program these into each experiment is, we believe, a major time-saver for the researcher who wishes to

concentrate on research, rather than technology.

Since its formal introduction eight months ago, *VisionWorks* has been a notable success. There are now 16 systems in the field, at prices ranging from $16,000 to $40,000. Considering that obtaining funds for such a large purchse typically takes almost a year, we think this very encouraging. We hope that in a few months, with the expiration of the first "grant cycle", we will se a considerable increase in orders.

BIBLIOGRAPHY

Coletta, N. J. and Williams, D. R. (1987) Psychophysical estimate of extra-foveal cone spacing. *Journal of the Optical Society of America,* A4, 1503-1513.

Derrington, A. M. & Lennie, P. (1984) Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *Journal of Physiology (London),* 357, 219-240.

Eizenman, M. Hallett, P. E. & Frecker, R. C. (1985) Power spectra for ocular drift and tremor. *Vision Research,* 25, 1635-1640.

Hirsch, J. and Miller, W. H. (1987) Does cone positional disorder limit resolution? *Journal of the Optical Society of America,* A4, 1481-1492.

Perry, V. Hugh & Cowey, Alan (1985) The ganglion cell and cone distributions in the monkey's retina: Implications for central magnification factors. *Vision Research,* 25, 1795-1810.

Smith, Robert A. & Swift, Dan (1990) Mapping the fine-grain sensitivity of the parafoveal visual field. *Investigative Ophthalmology and Visual Science,* 31(4), 495.

Smith, Robert A. & Swift, Dan (1991) Summation areas for equiluminant spots. *Investigative Ophthalmology and Visual Science,* 32(4), 1273.

Zrenner, E., Abramov, I., Akita, M., Cowey, A., Livingstone, M. & Valberg, A. (1969) Color perceptionL Retina to cortex. In Spillman & Werner (Eds.), *Visual Percpetion: The Neurophysiological Foundations.*

APPENDIX 1 StimulusMaker(tm) Functionality

It is not possible to explain the full functionality of StimulusMaker in a document of this size. The following is a simple list of the options on the various StimulusMaker menus, with some minimal commentary. This provides at least a suggestion of the level of functionality afforded. All stimuli are considered to be objects and as such share a common set of specifications. Specifications involving visual angle are based on the standard screen size of 28 by 21 cm, the standard mode of 1024 by 512 pixels, and the standard viewing distance of 75 cm. Additional specifications for each type of stimulus are listed in the following sections.

Feature	*Specification*
Color Mode	Standard RGB, Gamma-corrected RGB, Gamma-corrected Monochrome
Standard RGB	256 intensity values for each color gun
Gamma_RGB	User specifies intensity value for each gun from 0 to 255, display shows the linearized value that *StimulusMaker* actually uses
Monochrome	User specifies intensity from 0 to 32757 (15 bits). *StimulusMaker* provides a unique linearized intensity for each specified value
Type of Object	Patches and Forms
Type of Patch	Gratings, DCGs, Derivatives of Gaussians, Random-Dot Stereograms, checker-boards, rotating vanes (see Sections below for more details).
Type of Form	dots, circles, lines, rectangles, spheroids (see Sections below for more details)
Eye	Left, right, both (Use this feature for dichoptic presentation)
Disparity	-255 - +255 pixels. Each pixel is 1.2 arc-minutes. (However see sub-pixel disparity resolution feature in Section d.)
Number of Objects	1 - 100 (This is a fairly arbitrary value; can be increased if required)
Size (*Patch* only)	User specifies rectangular area in terms of pixels
Horizontal	1 - 1024 pixels in standard mode; each pixel is 1.2 arc-minutes

22

Vertical	1 - 512 pixels in standard mode; each pixel is 1.9 arc-minutes
Position	Center of patch can be at any pixel location, constrained only by size of patch.
Spatial Window (Patch only)	None, Vertical only, Horizontal only, both. User specifies standard deviation of Gaussian in degrees (floating point). Certain stimuli (e.g. random-dot stereograms, gratings moved with palette tricks) can not be spatially windowed.
Contrast Mode	Same_as, Constant, Windowed, Flicker, User-specified
Same_as	Any object can share the same Contrast profile as another within the same temporal interval. Aside from the added convenience of having to specify only one contrast profile, this feature allows efficient use of the available LUT slots, since the same LUT slots will be used by both objects.
Constant	Contrast is determined by the specified peak and trough brightness values and any additional contrast settings that can be specified on the stimulus menus
Windowed	Gaussian (User specifies peak and standard deviation in frames) Linear (Rise and Fall onset and offset specified in frames)
Flicker	Square-wave on-off, Square-wave counter-phase, Half-(rectified) sine-wave, counterphase sine-wave, raised cosine. The user can specify duty cycle and starting phase.
Frequency	Specified in cycles per second. This is a floating point value.
User-specified	User creates an ASCII file of contrast values; if length of file is less than duration of interval, values will repeat cyclically
Duration Mode	Fixed (User specifies number of frames) Until response (Key-press terminates stimulus).
Number of LUT slots	0 - 255. The total number of LUT slots used by all objects within a temporal interval can not exceed 255 (but see same_as feature above)
Tone Signal	A tone can be sounded coincident with visual stimulus onset. Duration is specified in number of frames, and frequency can vary from 1 to 14,000 Hz.

Trigger Pulse	A trigger pulse can be put out the parallel port coincident with stimulus onset. Polarity, channel, duration (in frames), and type (single or continuous pulse) are specifiable. A pulse can be initiated during any specified portion of an object. Pulses can also be initiated synchronously with flickering stimuli at any phase of the flicker. A secondary pulse can also be specified with flickering stimuli.

a. Specifications for Gratings

Feature	Specification
Types	Normal (i.e. linear, varying in one dimension) and Circularly Symmetric
# of components	1 - 99. Hardware will impose additional constraints if components have non-zero velocities.
Spatial frequency	.01 to 999 cpd (real upper limit depends on screen size and viewing distance: ~12 cpd for a vertical grating under standard conditions)
Phase	fraction of complete cycle from 0 to 1, where 0 (and 1) represent sine phase starting at left of patch
Angle	-90 - +90 degrees
Contrast	0 - 1
Velocity	-99 - +99 deg/sec. With standard frame rate of 120 Hz, real upper maximum is Nyquist limit of 60 cycles/second.
Motion methods	Blits, Palette Trick
Blits	pre-computes specified number of phases and blits them in at appropriate interval. The upper limit on patch size is about 69,000 pixels (about 10 by 10 degrees) updated every frame (138,000 updated every 2 frames, etc.) for a grating with one moving component. The limit reduces by a factor of 2 for each additional moving component.
# of Phases	1 to 999. Real upper limit depends on available D-RAM (memory). It is about 25 for maximum-size patch.

Palette trick
achieved by writing stimulus to the look-up table (LUT), and writing a sawtooth to the screen. There are no size limitations. However the use of palette tricks is restricted to components that share the same orientation and velocity. Furthermore, no spatial window is permitted.

Epsilon
0 to 1. Maximum allowed error (proportion) between intended and actual spatial frequency. Basic limitation is that a complete cycle of the grating must be completed in an integral number of pixels.

Minimize
With palette trick motion, the user can choose either to minimize the difference between the actual and specified spatial frequency, or to minimize the number of lut slots while still not exceeding epsilon.

b. Specifications for Differences of Gaussians (DOGs)

Feature	Specification
Dimensions	one or two (radially symmetric)
Sigma	0 to 999 degrees
Space Constant Ratio	.01 to 999
Phase	Standard or reversed
Angle	-90 to +90 degrees
Contrast	0 to 1
Contrast ratio	0 to 9 (the user can explicitly set the relative contrasts of the two gaussians)
Maintain equal area	*StimulusMaker* allows an option to automatically set the contrast ratio to maintain equal area above and below the zero-crossings.

c. Specifications for Derivatives of Gaussians (DRGs)

Feature	Specification
Derivative	0 to 99

25

Peak frequency	.01 to 999
Phase	Standard or reversed
Angle	-90 to +90 degrees
Contrast	0 to 1

d. **Specifications for Random-Dot Stereograms**

Feature	Specification
Shape	rectangular for both background and center area
Mode	static, dynamic or partial
Static	Basic mode that permits sub-pixel disparities
Dynamic	Random elements are modified every lut update (can be every frame if desired)
Partial	(-100 to +100) user specifies the degree of correlation between points in the center rectangles.
Dot Density	0 to 1 (default is .5). refers to proportion area covered by the random element that is assigned to the *crest* color
Dot Size	any number of pixels for both X and Y; minimum size is 1.2 X 1.9 arc-minutes under standard conditions
Center rectangle	*Size*: limited only by size of patch *Location*: anywhere within patch, constrained by whole-pixel boundary and size *Disparity*: resolution is 1.2 arc-minutes in standard mode; 0.28 arc-seconds in sub-pixel mode with static gratings

e. **Specifications for Checker-board Pattern**

Feature	Specification
Width of inner square	1 to 1024 pixels
Height of inner square	1 to 512 pixels

f. Specifications for Rotating Vanes

Feature	Specification
# of vanes	1 to 99
Type	Solid or Sinusoidal
Percent Blur of edges	0 to 100%
Inner radius	1 pixel minimum to outer radius maximum
Outer radius	Minimum: inner radius; Maximum: 512 y pixels
Cycles/sec	-66 to 66 (aliasing occurs at around 20 cps)

g. Specifications for Dot

Feature	Specification
Anti-aliasing	disabled or enabled (anti-aliasing allows placement of spot on a fractional pixel position)

h. Specifications for Circle

Feature	Specification
Anti-aliasing	disabled or enabled (anti-aliasing allows placement of circle on a fractional pixel position and fractional radius)
Radius	1 to 256 y-pixels

i. Specifications for Line

Feature	Specification
Anti-aliasing	disabled or enabled (anti-aliasing allows placement of

line on a fractional pixel position and fractional length)

Length   full screen, restricted only by location

j. Specifications for Rectangle

Feature	Specification
Anti-aliasing	disabled or enabled (anti-aliasing allows placement of rectangle on a fractional pixel position and fractional size)
Length	full screen, restricted only by location
Height	full screen, restricted only by location
Angle	-90° to 90°

k. Specifications for Spheroid

Feature	Specification
Types	(a) Raised Cosines and (b) Circularly Symmetric Gaussians
Radius	can be specified in degrees, mm, or pixels, but can not be less than 2 y-pixels nor greater than 256 y-pixels.

APPENDIX 2. VisionWorks(tm) Information

     The following are a sample of informational flyers concerning VisionWorks and its component subsystems. While scarcely complete, these offer some indication of the scope of the system.

VisionWorks™

A high performance 80486 PC-based stereoscopic computer graphics display system, stimulus generator and experiment controller designed for vision research. Modular software design allows integration of user written code and utility modules.

POWERFUL HARDWARE

- Binocular, monocular, dichoptic and stereo display formats using a TI TMS34020 graphics controller
- 120 Hz frame-rate
- RGB monitor with 256 colors out of a palette of 16.7 million (8bits/color gun)
- Special RGB to monochrome adapter that gives 15 bits of brightness resolution in monochrome
- Built-in gamma correction
- Stereo images presented at 120 Hz, with alternating left- and right-eye images, yield a flicker-free 60 Hz frame-rate to each eye.
- Trial frames may be worn under the goggles
- Master or slave interface to exts.nal devices
- Extremely high bit-image transfer rates and full double-buffering (4 display buffers plus scratchpad image memory) allow smooth flicker-free stereo animation at 1024 x 512 pixel resolution.
- Separate monitors for operator's console and graphics display
- Trackball for graphical input with an additional potentiometer for the Z axis
- 80486 (33MHz) cache computer with 100 MByte hard-drive and coprocessor

UTILITY MODULES

UNIQUE STIMULUS GENERATION

- Hierarchical menu-driven intelligent system which
  - allocates and optimizes internal resources
  - disallows illegal stimulus configurations
  - sets parameter values automatically with complete over-ride capability
- Two-dimensional Gaussian spatial window
- Many types of flicker
- Linear, Gaussian or user-input temporal window
- Powerful pre-defined stimulus selection, including
  - simple forms (bars, circles, etc.)
  - gratings with any number of components; each component with any spatial frequency, orientation, contrast, velocity, or phase
  - derivatives of Gaussian and differences of Gaussian
  - Random Dot Stereograms (RDS): static, dynamic, partially correlated

- New RDS mode yields disparity resolution exceeding 1 arc-second at normal viewing distances
- Sequencer allows multiple stimuli to be present at one time

The VisionWorks™ Photometric Linearization System

The VisionWorks Linearization System builds the system lookup tables that compensate for non-linearities in the color output DACs and (most importantly) in the monitor. Monitors have built-in "gamma-correction" which produces a non-linear output that compensates for the compressive non-linear brightness response of the human visual system. Since vision scientists wish to specify luminance in linear units, VisionWorks allows the user/programmer to specify output brightness in linear steps between 0 and the maximum luminance that the system has been calibrated to produce. In gamma-rgb mode there are 256 steps, in gamma-monochrome mode there are 32,768. Although each VisionWorks system is individually calibrated before being shipped, users who have critical requirements for accurate stimulus contrast, or who use different display formats or brightness may wish to be able to recalibrate systems themselves. Linearization cannot be performed using a standard photometer, although linearity may be checked with such equipment. The Linearization system works ONLY with VisionWorks systems.

The Linearization System is comprised of both hardware and software. The hardware includes a 12 bit A/D which fits inside your VisionWorks PC, a power supply and amplifiers, and a photodetector head containing a silicon photodiode and photometric filter. The amplifiers remain connected to the R, G and B channels and to the output of the grayscale expander at all times. The software is entirely menu driven and includes a variety of procedures, all simple to perform. These allow 1) setting the monitor's luminance range 2) building the gamma correction tables 3) checking linearity resulting from the tables, 4) measuring white field uniformity, and 5) measuring monitor stability.

The Linearization System works by measuring the output voltages of each DAC bit, and then measuring the precise voltage/brightness response of the monitor. This relationship is measured over 256 points in for each color channel and in the monochrome mode, which uses the grayscale expander. In monochrome mode, an interpolative smoothing algorithm is used to produce a smooth curve of 32768 luminance values between the measured data points. Each luminance value results from a voltage. A search algorithm establishes the closest voltages which the DAC/grayscale expander can produce for each of the 32,768 table entries. The R, G, B values which produce linear luminance steps are then built into lookup tables, which are loaded into GSP memory when the system is booted.

The Linearization system may be used to perform major or minor calibrations. A major calibration tests all components, including each RAMDAC bit, and takes about 30 minutes. Minor calibrations may test only a single component (such as the monitor) and need not actually rebuild the lookup tables. A typical calibration check will take less than 5 minutes.

The Linearization system is recommended for a variety of situations. These include: 1) when the monitor will be heavily used (causing phosphor changes), 2) when very accurate contrasts are required, and 3) when the user wishes to change the monitor image size or brightness settings.

The VisionWorks™ Experiment Workstation
developed and sold by Vision Research Graphics, inc.

What is VisionWorks?

VisionWorks is a PC-based, stereo-capable, computer graphics system specialized for use as a visual stimulator and experiment controller. VisionWorks is the first display system which allows a vision researcher to integrate menu-driven stimulus generation and psychophysical methods. This can be done with great flexibility, allowing creation of an unlimited number of new and unique experiments, with minimal programming.
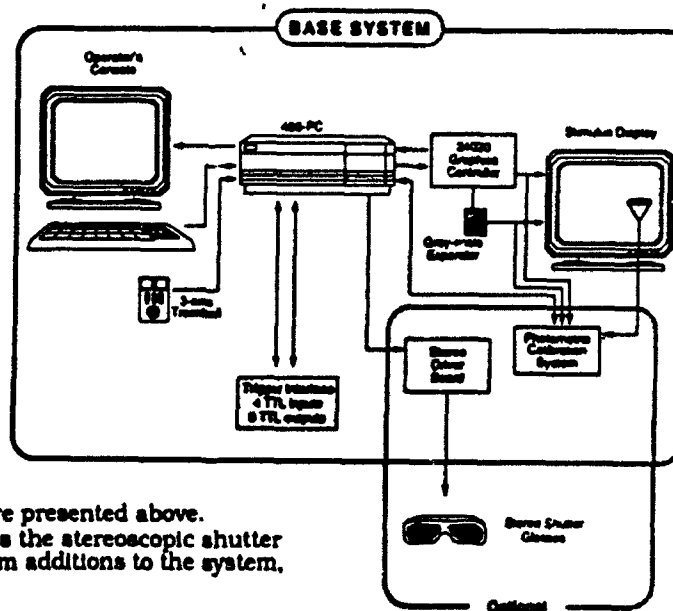
VisionWorks Components

VisionWorks is sold as a base system with hardware and software options that are priced separately.

Hardware

The base system hardware consists principally of a high-end PC and a modified high-resolution graphics controller and monitor. The base system has features indispensable to vision researchers, including a 120Hz frame-rate, 15-bit gray-scale resolution, and careful photometric calibration. VRG's optional stereo driver board and shutter glasses give the system stereoscopic capability.

VisionWorks™
Hardware

VisionWorks' hardware components are presented above.
In addition to standard options such as the stereoscopic shutter glasses, purchasers may request custom additions to the system, such as a read/write optical drive.
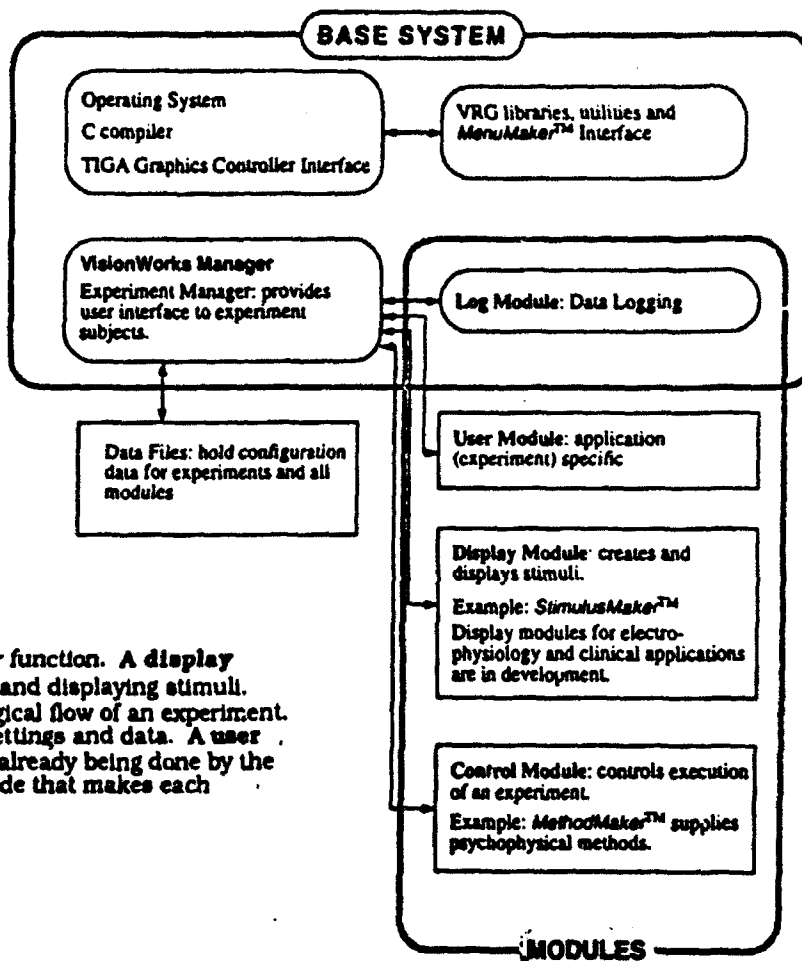
Software

The base system provides a complete software development environment,
which includes extensive function libraries, tools, and utilities useful to
vision researchers. VisionWorks' software organization is presented below.

Application programs execute
the experiments. Applications
are written using the
development environment.
An Application integrates
the VisionWorks Experiment
Manager with modules.

The Manager has two major
functions: 1) it provides the
matrix into which four types
of modules are integrated to
produce a complete application;
2) it provides the interface
for running and configuring
experiments. Experiments will
generally be run using parameters
stored in configuration files.

Each module performs a particular function. A **display
module** is responsible for creating and displaying stimuli.
A **control module** manages the logical flow of an experiment.
A **log module** records parameter settings and data. A **user
module** is used to do anything not already being done by the
other modules, and contains the code that makes each
experiment unique.

BASE SYSTEM

Operating System
C compiler
TIGA Graphics Controller Interface

VRG libraries, utilities and
MenuMaker™ Interface

VisionWorks Manager
Experiment Manager: provides
user interface to experiment
subjects.

Log Module: Data Logging

Data Files: hold configuration
data for experiments and all
modules

User Module: application
(experiment) specific

Display Module: creates and
displays stimuli.
Example: *StimulusMaker*™
Display modules for electro-
physiology and clinical applications
are in development.

Control Module: controls execution
of an experiment.
Example: *MethodMaker*™ supplies
psychophysical methods.

MODULES

Modules may be written by a user or purchased from Vision Research
Graphics (VRG). Modules currently available from VRG are *StimulusMaker*, a
display module which provides preprogrammed stimuli and *MethodMaker*, a
control module which provides psychophysical methods. The VRG modules
are large and powerful; a user would typically write small, special-purpose
modules to perform tasks that the VRG modules do not provide.

StimulusMaker™ and *MethodMaker*™

The *StimulusMaker* and *MethodMaker* modules use menus to allow the user to set up and view stimuli and psychophysical methods. However, some programming is needed to integrate these modules with the VisionWorks Manager and to allow *MethodMaker* to manipulate variables used by *StimulusMaker*. Applications that use these modules can be created with relatively small amounts of additional code, much of which can be copied from our examples. The code will comprise the User Module.

Purchase of both *StimulusMaker* and *MethodMaker* includes 8 hours of custom programming by VRG, typically enough to produce a simple working application and provide a good starting point for creating related experiments. Please see the next page for VRG's custom programming policies.

What can you do with VisionWorks?

VisionWorks provides an extremely flexible platform for developing and running experiments as a stand-alone psychophysical workstation, or in conjunction with external equipment for electrophysiology. The *StimulusMaker* and *MethodMaker* modules allow rapid development of experiments requiring spatially and temporally modulated stimuli, random dot stereograms, and standard psychophysical methods.

Stereoscopic capability makes VisionWorks ideal for investigating and simulating strabismus and amblyopia, or developing clinical screening procedures. Used with a large monitor, VisionWorks is highly suitable for developing advanced perimetry techniques which use chromatic and spatially and temporally modulated stimuli. Anti-aliasing techniques create effective sub-pixel resolution for vernier acuity and stereograms.

Examples of Possible VisionWorks Applications

- Electrophysiological and psychophysical investigation and simulation of vision deficits in strabismus and amblyopia.
- Electrophysiological and psychophysical response to spatio-temporal and novel stimuli for investigation of visual field deficits.
- Sophisticated automated vision testing using spatial and RDS stimuli.
- Comparisons of stimuli generated using luminance and chromatic contrast.
- Visual Field (scotoma) mapping using varying spatio-temporal target and background parameters.

VisionWorks Custom Programming Policies

If you purchase VisionWorks directly from VRG, with both the StimulusMaker and MethodMaker modules, you are entitled to 8 hours of free programming. This amount of time is enough to produce a simple VisionWorks application which uses a subset of the stimuli available in StimulusMaker, and a subset of the psychophysical methods available in MethodMaker. The principal purpose of this simple application is to provide a programming example upon which you can model your own code.

Your free application should not require any changes in the VisionWorks base system, StimulusMaker, or MethodMaker, unless the changes are already planned. If your program requires planned base system changes, there will be no additional charge but the completion date for the program can not be guaranteed. If the program requires changes in the base system which were not planned, VRG will provide the programming at a negotiated price and schedule.

If the specifications for your program have been received and approved by VRG by the time the PO is placed, the program will be shipped with the VisionWorks system. If the specifications are provided after receipt of the PO, the system may be shipped without the program. Payment for the system is due 30 days after shipment, irrespective of inclusion of the application program.

VRG will supply custom programming on an hourly or on a contractual basis. The hourly rate is $75/hr. and there is an eight-hour minimum for contracted jobs. Payment for custom programming is due as follows: 50% upon receipt of PO, the remainder payable on a mutually-agreed-upon schedule.

3/92

APPENDIX 3. pcStereoscope(tm) Information

The following pages contain a sample of informational flyers concerning pcStereoscope. These offer further insight into to scope and usefulness of this project.

pcSTEREOSCOPE™
SOFTWARE DEVELOPMENT SYSTEM

A programming environment specialized for stereoscopic 3-D

The pcSTEREOSCOPE Software Development System speeds development of stereoscopic 3-D programs by providing high-speed graphics functions, animation and menuing.

Stereoscopic graphics applications include:

- Vision Research
- Presentation of multidimensional data
- Scientific visualization and biomedical imaging

- Fractals
- Geologic imaging
- Games
- CAD/CAM

The Development System includes the following categories of routines:

Random-Dot Stereogram
Angle Conversion
Angle, Speed, and Motion
BitBLT
Sprite
Geometric Primitives,
   including arc, line, clip,
   filled polygon
Hercules Monitor (provides
   separate monitors for
   text and graphics)
Math
Palette

Screen Conversion
Stereo
Text
Time Conversion
Track Ball

Utilities:
Browse Module
Menu Module
Object Coloring Module
Story Module for
   animation/image
   sequencing

Vision Research Applications

The *Software Development System* allows the vision researcher to develop his/her own experiments and demonstrations. Random dot stereograms are extensively supported in the system. Applications for vision research include:

- stereoscopic and dichoptic illusions
- moving wire frame figures, with or without stereo or perspective cues
- aftereffect experiments (color, motion, grating)
- static, dynamic and moving texture patterns and random dot stereograms

Software Development System

Requires pcSTEREOSCOPE Hardware and Borland Turbo C Version 2.0. Includes C Source Code examples, User's Manual, and Programmers' Reference Manual.

Cost $350

pcSTEREOSCOPE Hardware

An inexpensive stereoscopic 3-D system for EGA or VGA equipped IBM PC's or compatibles.

The pcSTEREOSCOPE time-multiplexes left-eye and right-eye images, yielding a 30 Hz frame rate at each eye to provide color 640 x 350 pixel images.

The hardware package consists of a single short-slot card, one pair of liquid-crystal shutter glasses and a connector. A trackball with an additional potentiometer for the z axis is available as an optional 3D input device.

Cost $450
Additional Shutter Glasses $75/pair
Trackball $150

Recommended System Configuration:

12 MHz PC AT compatible or better, with hard drive, 1 MB minimum memory, and VGA graphics.

VISION RESEARCH GRAPHICS inc.

99 Madbury Road
Durham, NH 03824

Phone: (603) 868-2270
Fax: (603) 868-1352

pcSTEREOSCOPE Software Development System
Technical Description

The pcSTEREOSCOPE Software Development System (PCSDS) is a C-language based development environment for the Vision Research Graphics, inc. pcSTEREOSCOPE hardware. The PCSDS consists of libraries which provide:

1. drivers for the hardware,
2. very fast graphics primitives,
3. advanced graphics operations,
4. sequencing/animation routines,
5. a menu/forms input system,
6. a browse/help utility.

The PCSDS is intended for use with Borland Turbo C 2.0. Extensive documentation and application source code examples are provided. An optional driver development kit contains driver source code suitable for incorporation into existing programs. Any C compiler can be used for driver development.

Converting Application Programs to Stereo
The PCSDS is suitable for both programming 'from scratch' or for converting existing "3-D" programs to stereo. Converting programs which generate static images (such as CAD programs) is generally simple. The program's existing routines are used to place left-eye and right-eye images into two display pages. The pcSTEREOSCOPE driver is then called, which turns on the shutter glasses and flips the display pages.

Display Formats
The PCSDS principly supports IBM mode 10H, which is the highest resolution standard IBM mode with two display pages. This mode provides two display pages of 640 x 350 pixels with 16 colors. The driver routines can easily be modified to work with any mode that supports two display pages, including super EGA/VGA. However, the graphics primitives will generally not work in modes beyond 800 x 600 x 4.

Graphics Primitives, Advanced Operations, and Utility Routines
The development system provides a wide variety of primitives which write directly to the graphics hardware, and are in general the fastest EGA/VGA primitives that we know of. These include points, lines, circles, elipses, polygons, and clipping. Advanced routines include bitBlit, virtual sprites, bitplane operations, palette manipulation, screen-to-disk and disk-to-screen copies. Many utility routines are supplied to simplify coordinate conversion, math, use of multiple monitors, angle conversion, etc. A trackball with an additional pot for Z is supported as a 3-D input device.

Technical Support and Programming Services
The development system comes with 30 days of support. Extended support packages are available at extra cost. Vision Research Graphics also provides custom programming services.

pcSTEREOSCOPE Technical Description

What is a Stereoscopic Display?

A stereoscopic display provides separate images to the left and right eyes. When these images simulate the slighly different viewpoints of our two eyes, the brain will combine the ima 's su:h that we see them with an actual sensation of depth. This differs greatly from so-called 3D displays, which use perspective cues, shading and superposition to convey merely an impression - not sensation - of depth. Using both schemes together provides the static depth cues available in real world viewing, and yields the strongest perception of depth.

Theory of Operation

The pcSTEREOSCOPE works by time-multiplexing left- and right-eye images. The images are alternated on a CRT monitor, while shutter glasses allow each eye to view only the images intended for that eye. Since a standard monitor runs at 60 Hz, the frame-rate to each eye is 30 Hz.

Hardware Description and Requirements

The pcSTEREOSCOPE consists of a pc board, liquid-crystal shutter glasses, and a connector. The pc board occupies a single 8-bit slot in an IBM PC/XT/AT or compatible, or in an IBM PS2 Model 30. The connector is placed between an EGA/VGA controller and monitor, and provides a line which carries the vertical sync signal to the pcSTEREOSCOPE board. The board detects the sync pulse and interrupts the computer. An interrupt driver then directs the board to open and close the shutters over the left and right eyes, and sets the frame buffer address to display the image to the proper eye.

The pcSTEREOSCOPE works with any EGA/VGA controller. Most monitors are suitable, exceptions being those with long-persistence phosphors. Typically IBM mode 10H is used to provide 640 x 350, 16-color graphics. Stereoscopic images are produced by generating left- and right-eye images in the two pages of the display, and then turning on the software driver. This mode is extensively supported in our pcSTEREOSCOPE Software Development System. Other display modes may be used, including super-EGA/VGA modes, so long as the graphics board provides two full display pages in the desired mode.

Vision Research Graphics, inc.
99 Madbury Rd.
Durham, NH 03824

Phone: (603)868-2270
Fax: (603)868-1352

Stereo-Snapshot/SVGA

A collection of (TSR) utilities that allow EGA, VGA, and Super-VGA images to be captured and used in stereo-pairs. After capturing images for the left and right eyes, the images may be viewed in stereo, using the pcSTEREOSCOPE hardware. Stereo images can be sequenced from storage on a hard drive.

The output of most 3D programs can be viewed in stereo by using this program. It works with 3D CAD, solid modeling, video frame capture, 3D graphing, data presentation, and many others. A public domain 3D CAD program is included to demonstrate Stereo-Snapshot/SVGA.

This product requires the pcSTEREOSCOPE hardware and a Super-VGA board containing the Tseng 4000 chipset.

Modes supported are:
640 x 480 x 16 color
640 x 480 x 256 color
800 x 600 x 16 color ----*Note:* The driver for this mode has been written and tested, but is not yet implemented. If you want this mode, there will be a delay after receipt of the PO.

Note: 800 x 600 x 256 color would probably work but has not been tried. We will try it on a custom programming basis.

Prices
Stereo-Snapshot/SVGA software $400
pcSTEREOSCOPE hardware $450

© 1992

6/92

VISION RESEARCH GRAPHICS, inc.

99 Madbury Road • Durham, New Hampshire 03824
(603) 868-2270 • FAX (603) 868-1352

PRICE LIST

pcSTEREOSCOPE™

HARDWARE: *When ordering, please specify EGA or VGA cable.*
Stereoscopic Computer Graphics Display System........................$500
Requires IBM PC-XT, PC-AT or Compatible with EGA/VGA card and monitor
> Package Includes:
> PC short-card display driver board (Supports 5 prs. of glasses)
> One pair liquid crystal shutter glasses
> One EGA or VGA adapter cable
> Hardware Installation Instructions

SOFTWARE: *Please specify disk size: 5.25" HD, 5.25" LD, or 3.5" HD*
pcSTEREOSCOPE HARDWARE required for all software.

The VisionLab™ $600
> Package Includes:
> Over 30 experiments and demonstrations
> Laboratory Manual/ User's Manual
> Single User License

VisionLab with a site license for 1 extra station $1100
VisionLab with a site license for unlimited stations $1500

pcSTEREOSCOPE Software Development System........................$350
Requires Turbo C. 2.0 or Turbo C++
> Package Includes:
> C Source Code examples
> Function Library, and driver source code
> User's Manual/ Programmers' Reference Manual
> Single User License

ADDITIONAL ITEMS PURCHASED SEPARATELY:
The VisionLab Manual $25
Development System Manual...35
EGA Adapter Cable..30
VGA Adapter Cable..35
Adaptor for Multiple Glasses..30
3-Axis Trackball 155

International Orders, Consignee pays shipping charges, duties and taxes.
Domestic Orders, Shipping and Handling, $15 for ≤ $1500, $25 for > $1500

MasterCard and VISA accepted. Prices are U.S. Currency and are subject to change.   1/93

II. PUBLICATIONS

Smith, Robert A. & Cass, Peter F. (1989) "Effect of eccentricity on spatial summation and acuity. *J. Opt. Soc. Amer. A*, 6 #10, 1633-1639.

Smith, Robert A. & Swift, Dan (1990) Mapping the fine-grain sensitivity of the parafoveal visual field. *Investigative Ophthalmology and Visual Science*, 31(4), 495.

Smith, Robert A. & Swift, Dan (1991) Summation areas for equiluminant spots. *Investigative Ophthalmology and Visual Science*, 32(4), 1273.

Smith, Robert A. & Cass, Peter F. (1992) Orientation and motion effects in neural aliasing. In review, *Journal of the Optical Society of America*.

III. PROFESSIONAL PERSONNEL

Robert A. Smith, PhD.
Dan J. Swift, PhD.

IV. PROFESSIONAL INTERACTIONS

Smith, Robert A. "Mapping the fine-grain sensitivity of the parafoveal retina." presented to ARVO, May 1990.

Swift, Dan J. "Encoding for stereoscopic depth determined from perceived disparity shift". presented to ARVO, May 1990.

Smith, Robert A. & Swift, Dan "Summation areas for equiluminant spots." to be presented to ARVO, May 1991.

Swift, Dan J., Smith, Robert A. & Panish, Steven "True three-dimensional capture and illusory contours." poster presentation at the 1990 Annual Meeting of the Optical Society of America, Boston.