

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A151 888

AFOSR-TR- 85 - 0248

2
WR

VISION ALGORITHMS AND PSYCHOPHYSICS
(Annual Report 1983-84)

Investigator
(617) 253-5776 or 5700
Arpanet: WHIT@MIT-OZ
Contract No. F49620-83-C-0135

Whitman A. Richards E10-120
Massachusetts Institute of Technology
Cambridge, MA 02139



ABSTRACT: Vision by man or machine is the construction of useful symbolic descriptions from images of the world. Studies of the human visual system provide valuable insights into the kinds of descriptions that will be the most useful, but little insight into the computational problems involved in deriving and manipulating these descriptions. This research examines several computational problems associated with aspects of two- and three-dimensional vision. The solution to these problems includes the design and implementation of particular algorithms. Their efficiency and flexibility is compared with that of the human visual processor.

Key words: Image understanding, visual pattern recognition, vision algorithms, human vision, biological information processing.

September 1984

OTIC FILE COPY

SELECTED
APR 1 1985
A

Approved for public release;
distribution unlimited.

85 03 13 20 6

CONTENTS:

- 1.0 Introduction: The Problem and Goal**
- 2.0 Development of Codon Theory**
 - 2.1 Transversality and "Parts"**
 - 2.2 Constraints on 2D Shapes**
- 3.0 Codon Implementation**
 - 3.1 The Algorithm**
 - 3.2 Sketch of the "See" Machine**
- 4.0 Groupings and "Glue"**
 - 4.1 Color**
 - 4.2 Texture**
- 5.0 Visual Motion**
- 6.0 Summary**
- 7.0 Appendix 1 - Software Developed**
 - Appendix 2 - Publications and References**

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFOSR-TR- 85-0248	2. GOVT ACCESSION NO. AD-A151888	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) VISION ALGORITHMS AND PSYCHOPHYSICS		5. TYPE OF REPORT & PERIOD COVERED Annual Report 15 Jul 83 - 14 Jul 84
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Dr. Whitman A. Richards		8. CONTRACT OR GRANT NUMBER(s) F49620-83-C-0135
9. PERFORMING ORGANIZATION NAME AND ADDRESS Dept of Psychology Massachusetts Institute of Technology Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2313/A5 61102F
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Office of Scientific Research/NL Bolling AFB, DC 20332		12. REPORT DATE September 1984
		13. NUMBER OF PAGES 19
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Image understanding, visual pattern recognition, visual algorithms, human vision, biological information processing.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Vision by man or machine is the construction of useful symbolic descriptions from images of the world. Studies of human visual system provide valuable insights into the kinds of descriptions that will be the most useful, but little insight into the computational problems involved in deriving and mani- pulating these descriptions. This research examines several computational problems associated with aspects of two- and three-dimensional vision. The solution to these problems includes the design and implementation of particular algorithms. Their efficiency and flexibility is compared with that of the		

Annual Scientific Report
F49620-83-C-1-0135
September 1984

VISION ALGORITHMS AND PSYCHOPHYSICS

Department of Psychology
Massachusetts Institute of Technology
Cambridge, MD 02139

Dr. Whitman A. Richards

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution _____	
Availability Codes	
Avail and/or	Special
Dist	Special
All	

DTIC
COPY
INSPECTED
1

Controlling Office: Air Force Office of Scientific Research/NL
Bolling Air Force Base, DC 20332

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH/NL
NOTICE: THIS REPORT IS AVAILABLE TO DTIC
This technical report is approved for distribution and is
distributed as follows: 12.
MATTIE
Chief, Technical Information Division

VISION ALGORITHMS AND PSYCHOPHYSICS

1. Introduction: The Problem and Goal

"Seeing" requires the construction of symbolic descriptions of the external world. The most useful symbolic descriptions will be representations for each of the various objects in the three-dimensional scene. These objects, in turn, may be broken down further into more detailed modular representations that may include the various attributes of each object such as its color, texture, or the shape and relative motion of its parts. These latter properties are thus our basic building blocks from which more complicated descriptions are built. Vision understanding requires showing how such object properties can be represented internally, and how they can be brought together to create a description suitable for recognition or manipulation. This then, is our goal: to propose and implement a scheme for representing 3D shapes in a manner suitable for recognition.

To reach this goal, the research is proceeding along several parallel tracks. The first is the development of a theory for representing 3D shapes, or their 2D projections onto the image. Such a theory requires that the shape be spatially isolated. Hence the second research track is the identification of object candidates, using visual motion (or stereo). A third track is a machine implementation of the proposed schemes. And finally, a fourth research area interwoven among the others is the psychophysical explorations that provide hints about viable vision algorithms.

2. Development of the Codon Theory

Shape is one of the most important ways of categorizing and identifying objects. In Figure 1, the very simple shape of an eye immediately implies "animal". Seeing the "beak" would further constrain the class to be "bird". In contrast an isolated patch of texture is almost meaningless. This simple observation shows that rather simple shapes can provide a powerful representation for recognizing objects. Silhouettes or cartoons reinforce this notion. What, then, are the basic elements from which we can build simple shapes and make such powerful inferences from image contours?

In 1982, Hoffman and Richards proposed a primitive representation for the shape of 2D or plane curves. The key concept was that the representation should make explicit the parts of a shape (or 3D object), because objects are described most naturally in terms of their "parts".




Figure 1 The left-hand panels show two portions of the bird—one a texture, the other a "shape". Clearly the simple shape of an eye alone provides an important pointer to the class of object, namely "animal", whereas the texture patch alone offers few clues.

To find the parts of an object or even of a plane curve, one notes that when two parts are joined, a concavity in the surface will be formed, which appears as a minima of curvature or cusp in the image (see Fig. 2). This concavity property is a transversal one—stable under perturbations of the way the parts may be joined. Transversality is thus a fundamental regularity of natural objects. It is the basis of our scheme for decomposing a curve into "parts" suitable for recognition.

To represent the shape of the "parts" isolated by the concavities, or minima of curvature, Hoffman and Richards (1982, 1983, 1984) propose as a first abstract description that we use the extrema of curvature. These are the maxima, minima and zeroes of curvature along a plane curve. Such a representation has the important feature that it is invariant under similarity transforms—rotation, translation or dilation. The basic elements of the representation are called "codons", which are illustrated in Fig. 3. Each codon simply represents one of five possible relations between the maxima, minima and zeroes of curvature. They are identified by their number of inflections (zeroes). Shapes described in terms of a sequence of codons have several interesting formal properties, such as making skewed symmetry explicit, and being sensitive to the choice of figure and ground (see Fig. 4)—two important perceptual attributes (Hoffman and Richards, 1982, 1984).

Over the past year, theoretical work on the codon scheme for representing shapes has focused on two problems: (a) a rigorous mathematical statement of the transversality motion




Figure 2 Joining parts generally produces concavities in the silhouette.




Figure 3 The primitive codon types. Zeroes of curvature are indicated by dots, minima by slashes. The straight line (∞) is a degenerate case included for completeness, although it is not treated in the text. (See Richards and Hoffman, 1984, for definitions.)

and its generalization to smooth shapes, and (b) topological constraints on possible smooth 2D (image) contours defined by codon strings.

First we will summarize some work on the transversality regularity that is critical to a formal definition of a "part boundary".

2.1 Transversality and "Parts"

When two surfaces intersect they intersect *transversally*. This means that the tangent planes to the two intersecting surfaces are of different orientations at each point where they intersect, implying that there is a discontinuity of the tangent plane to the surface of the




Figure 4 Skewed symmetry is obvious in the codon string because half the sequence is reversed, ignoring the sign of the codon (left frame). Figure-ground reversal changes the codon string because maxima and minima of curvature are exchanged, providing a simple explanation for Rubin's face-vase illusion.

new composite object at each point along the contour of intersection (see Fig. 2). Contours of concave discontinuity are thus the part boundaries.

Consider now what happens if the concave discontinuity at the part boundary is smoothed as if a membrane were stretched across the discontinuity in the intersecting surfaces. Where then is the part boundary on the smoothed surface? Intuitively, one might choose the locus having greatest curvature. Hence we have proposed the following rule for partitioning smooth surfaces into parts:

Negative Minima Partitioning Rule: Divide a surface into parts at negative minima of the principal curvatures along their associated lines of curvature.

The proof that this rule indeed captures our intuitive notion of the part boundary is quite difficult, but has recently been completed by Bennett and Hoffman (1985). The thrust of the proof is to show that the smoothing of a concave discontinuity on a surface will produce local extremum of surface curvature in the neighborhood. The proof thus provides a solid mathematical foundation to our part boundary notion for 3D shapes. Given this mathematical rigor, we can now determine how the 3D boundary will appear when projected into the 2D image. Although the boundary will generally *not* be a point of extremal curvature, we expect that 2D extrema of (negative) curvature will lie near the projection of a 3D part

LEGAL CODON PAIRS		THIRD CODON				
		0 ⁻ - -	0 ⁺ + +	1 ⁻ - +	1 ⁺ + -	2 - -
TAIL HEAD		TAIL HEAD	TAIL HEAD	TAIL HEAD	TAIL HEAD	TAIL HEAD
0 ⁻ 0 ⁻	- -	[+ +]		+		[+ +]
0 ⁻ 1 ⁻	- +	+	+		[+ +]	+
0 ⁻ 2	- -	[+ +]		+	+	[+ +]
0 ⁺ 0 ⁺	+ +		[+ +]	+	+	
0 ⁺ 1 ⁺	+ -	+	+	[+ +]		+
1 ⁻ 0 ⁺	- +	+	+		[+ +]	+
1 ⁻ 1 ⁺	- -	[+ +]		+	+	[+ +]
1 ⁺ 0 ⁻	+ -	+	+	[+ +]		+
1 ⁺ 1 ⁻	+ +		[+ +]	+	+	
1 ⁺ 2	+ -	+	+	[+ +]		+
2 0 ⁻	- -	[+ +]		+	+	[+ +]
2 1 ⁻	- +	+	+		[+ +]	+
2 2	- -	[+ +]		+	+	[+ +]
NUMBER OF LEGAL PAIR SUBSTITUTIONS		5	2	3	3	5
NUMBER OF PAIR SUBSTITUTIONS		9	4	6	6	9

Figure 5 Table 1. Legal smooth codon triplets. The third codon can either follow or precede the pair. A (+) indicates a proper join. Because of symmetry, there are an equal number of total pluses in the head and tail columns.

boundary. The most common case where the projection has a cusp in the occluding contour has already been treated by Hoffman and Richards (1984).

2.2 Constraints on Codons

Clearly any plane curve can be described by a sequence of codons, for all curves can be characterized by a sequence of the extrema of curvature. However, once one imposes restrictions on the behavior of a plane curve, the sequence of curvature extrema may not be arbitrary. One example is the class of plane curves that are smooth and have no cusps. Included in this class are the smooth plane curves that represent the canonical outlines of smooth 3D shapes.

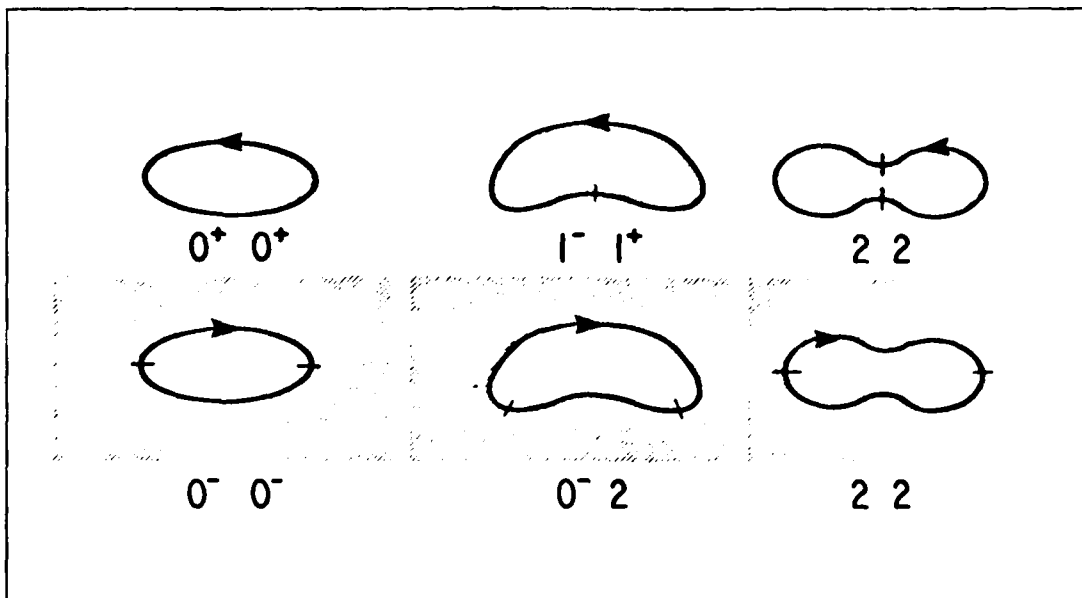


Figure 6 Legal smooth, closed codon pairs. Figure is indicated by cross hatching. Part boundaries are noted by the slashes. The dots indicate positive minima, which are used for shape descriptions but are not part boundaries.

To see that not all sequences of codons are possible if the curve is smooth, refer to Fig. 3 once again. Note that a 1^- can not follow a 1^- codon unless a cusp is allowed. Similarly, a 1^+ can not follow a 1^+ , because if such a join is attempted either a cusp will be created or, if the curve is indeed smooth, the 1^+ codon would have to be transformed into a type 2. To specify all legal smooth codon strings, we will first enumerate all pairs, and then show what pair substitutions are legal for one element in a sequence of pairs, thereby creating all possible triples.

Define the "tail" of a codon as the region about the first minima encountered when traversing the curve. The "head" of the codon is the subsequent minima. A smooth string of two codons is then allowable only if the head of the first codon has the same sign of curvature as the tail of the second codon in the string. To enumerate the possible codon pairs for a smooth contour, we require that the curvature of both the head and tail of a middle codon match the tail of its successor or the head of its predecessor in the string. All such legal pairs are given in the left column of Fig. 5. There are only 13 legal pairs out of a possible 25 combinations.

If we now require that the curve be closed smoothly on itself, then this constraint drastically reduces the number of legal pairs, for now the head and tail of the pairs must have the same sign. Inspecting Fig. 5, we see immediately that only $0^- 0^-$, $0^- 2$, $0^+ 0^+$, $1^- 1^+$ and $2 2$ qualify. These shapes are shown in Fig. 6. Surprisingly, now there are only three such legal shapes out of the possible 25 combinations! According to codon theory,




Figure 7 Legal smooth, closed codon triple and quadruples.

the "ellipse", the "peanut" and the "dumbbell" are the three most primitive shapes. Note that the ellipse has no "parts", the peanut has one part boundary, and the dumbbell has two parts. It will be these three primitive shapes that our implementation will seek in images, to be described in a later section.

In a similar manner, we can enumerate the class of all possible smooth-closed shapes that are topologically similar in their "bumps" and "dents". There are not very many forms of such curves. For example, a seven codon sequence has 78,125 possible combinations, but only 65 are allowable (see Richards and Hoffman, 1984, for proof). Fig. 7 shows legal smooth, closed plane curves for codon strings of length three and four. Note that even four codon elements can yield quite descriptive-looking shapes, such as the "fetus" or "animal" in the lower right. Furthermore, it should be obvious that combinations of closed codon shapes embedded within one another can represent a wide range of complex figures. The "eye" of Fig. 1, which is simply one ellipse within another, is one example. Or a "face" which is often depicted by an ovoid with two simple ellipses for the "eyes" and a "peanut"




Figure 8 A Gaussian pyramid structure using the mask shown on the left produces a "pyramid" of images (also shown on the left). These images are used to obtain the two binary pyramids of image "Snoopy", as shown in the right panel.

shape for the mouth, would be another example. How, then, can the codon shapes be extracted from images in order to build such representation?

3.0 Codon Implementation

3.1 Description of Algorithm

In order to represent codons within codons, such as in the eye or face example, we Gaussian filter our images at several scales, using a pyramid scheme proposed by Burt and Adelson (Burt & Adelson, 1983; Burt, 1982). Figure 8 shows the output of this first stage of processing of the image "Snoopy". Note that the algorithm gives us two pyramids—one capturing the "dark" blobs, the other the "light" blobs. We now are able to create a blob hierarchy, where blobs within blobs are specified as a linked-list tree-structure. Because Gaussian masks are used, we are guaranteed that our blob hierarchy will be well behaved (Yuille and Poggio, 1984; Koenderink, 1984).

With all the blobs located (some in more than one pyramid level), we next generate an edge list for each blob. Starting at the top of the blob, we encode the edge in a counter-clockwise fashion. Our algorithm is an adaptation of a standard edge crawl, using 8-way connectivity. Between a pixel and the previous pixel in the edge list there is a "tangent" — one of eight directions. To find the next pixel in the sequence we project a vector normal to this simple tangent vector (90 degrees clockwise rotation) and then sweep this vector counter-clockwise until it finds the next pixel. In this fashion we "hunt" between object and

background and generate the edge list. The implementation is simplified by using only the eight possible vectors in an 8-connected area around the pixel.

At this point a tangent could be computed at each edge point using a standard set of edge masks. Since we are dealing with "blobs" it makes more sense to compute the principal normal. This also turns out to be computationally efficient. Whereas the tangent points along the edge of the blob, the principal normal points towards the center. Consider then an edge point and the points immediately surrounding it. Each surrounding point that is part of the blob (signal) can be thought to pull on rubber bands attached to an imaginary vector emanating from the edge point. The sum of these pulls will tend to point the vector towards the center of the blob, and hence the vector will approximate the principal normal. The computational scheme for implementing this algorithm is described in Dawson and Treese (1984). The output of the algorithm is thus the normal to the outline of the blob. The tangent at each point is simply 90° to this orientation. Figure 9 shows one example for the "dumbbell"-shaped eyebrow of Snoopy.

Because the codon scheme is based upon extrema of curvature, we now must differentiate the tangent versus arc length along the blob outline. In the upper right panel of Figure 9 the tangent versus arc length is given by the upper curve of the graph. Its derivative is obtained simply by applying the "edge" operator shown in the same panel to obtain curvature versus arc length (lower curve on graph). The extremities of curvature used to specify the codons and hence the shape of the blob can now be read off directly (see Dawson and Treese, 1984).

3.2 Sketch of the "See" Machine

A large part of the effort over the past year was devoted to software and hardware development. Our "See" machine is a VAX 11/750 computer with a 400 megabytes fixed disc, linked to an Adage 3000 image processor. The Adage has 512 x 512 x 24 bits resolution and is our principal image processing device, performing the Gaussian pyramid convolutions in about 17 seconds. It is also used to generate "natural" images. Other peripherals include several single-frame graphics terminals, a matrix color camera, a Fairchild CCD camera and several vidicons eventually intended for color and stereo input. Still under development is an Ethernet connection to the Artificial Intelligence lab, and the capability for inputting motion sequences taken with a portable video camera.

The VAX runs Unix 4.2. Over the past year much special purpose software has been written to make the "See" Machine a user-friendly system suitable for both graphics and




Figure 9 The upper left panel illustrates the principal normal encoding scheme for the sub-blob of Snoopy's eyebrow. The upper right shows how curvature along the blob's outline is computed. A second example for the super ordinate blob of Snoopy's head outline is shown in the third panel.

implementation of the codon scheme. Over one-half man-years has been spent solely on software development (see Appendix 1 for a list of packages written).

4.0 Groupings and "Glue"

Most attempts to interpret images focus upon image contours. Hence a large part of image processing is concerned with edge detection algorithms. Our early attempts to isolate shapes also proceeded in this manner, using an algorithm called "cartoon" (Richards et al., 1982). Such edge-finding schemes when used on natural images are almost guaranteed to yield broken contours, which must be subsequently "glued" together appropriately. One method for identifying the pieces of contour that should be linked is to provide color or

texture labels.

Our current scheme which is based upon "blobs" does not suffer this problem, for all "blobs" are guaranteed to have closed boundaries. Nevertheless, sometimes different parts of the same object will appear as isolated blobs (such as the two eyes in a face or during occlusions) and it is useful to be able to assign material-property labels to the isolated blobs to provide indices for appropriate groupings. Color, texture and motion are the prime candidates for such labels.

4.1 Using Spectral Information to Represent Material Categories.

Earlier, Rubin and Richards (1982) had shown how an operator called the spectral cross-point could be used to find material changes across an edge. This condition is depicted in Fig. 10. Also shown in the same figure is a second condition, called the opposite-slope sign, which can be used to categorize the spectral composition of surfaces. In brief, this new condition describes the crude spectral shape of a pigment in terms of its derivatives of absorption versus wavelength. Details are given in a recent AI Memo 764 by Rubin and Richards (1984). The theory has implications for both psychology and neurophysiology. In particular, Hering's notion of opponent colors and psychologically unique primaries, and Land's results in two-color projection can be interpreted as different aspects of the visual system's goal of categorizing materials. Also, the theory provides two basic interpretations of the function of double-opponent color cells described by neurophysiologists.

4.2 Texture Fields

We have just begun some work on representing the local texture properties of a "blob" in terms of the four components of a linear flow field (i.e., dilation, rotation, shear and deformation). This approach is new because it attempts to infer the local organizations of the texture directly without first establishing correspondence (Richards, 1984). There are some formal similarities to Ullman's work in structure-from-motion which are also being explored.

5.0 Visual Motion

Being able to compute the movement in the changing retinal image is critical to an implementation of the codon theory, for it is perhaps the single most powerful method of isolating "blobs" of greatest interest. Such a description of the movement is not provided

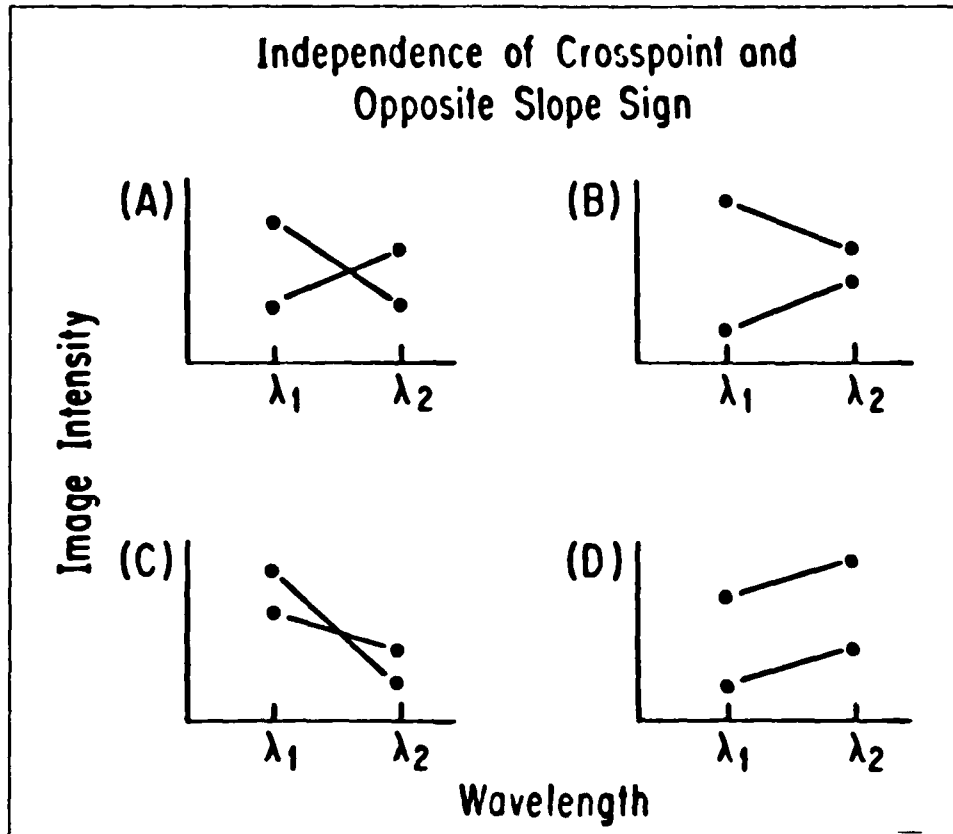


Figure 10 Graphs of image intensity (ordinate) versus wavelength (abscissa). Two wavelength samples, λ_1 and λ_2 , are shown. An image region yields two samples of intensity, one for each wavelength, and is represented by the line segment connecting the two sample values. a) & c) Two examples of the spectral crosspoint (Rubin & Richards, 1982). a) & b) Two examples of the opposite slope sign condition. This is the minimal configuration that shows different ordinalities. Note that the crosspoint and opposite slope sign condition are completely independent, since they can occur together (a), or each can occur alone (b and c), or neither can occur (d).




Figure 11 Ambiguity of the velocity field. (a) the arrows represent two possible velocity fields that are consistent with the changing image. (b) The curve C_1 rotates, translates and deforms over time to yield the curve C_2 . The velocity of the point p is ambiguous.

to our visual system directly, however; it must be inferred from the pattern of changing intensity that reaches the eye. Hildreth (1984a,b) has studied this problem in detail.

One serious obstacle to computing general motion is that the local motion measure-




Figure 12 (a) **Left:** Rotating spiral. (a) The true velocity field for a logarithmic spiral rotating in the image about the point O . (b) the initial perpendicular velocity vectors. (c) The computed velocity field of least variation. (b) **Right:** The barberpole illusion. (a) A circular helix on an imaginary cylinder, rotating about the vertical axis of the cylinder. (b) The two-dimensional projection of the helix and its velocity field. (c) The initial perpendicular velocity vectors. (d) The computed velocity field of least variation.

ments generally provide only one component of the local velocity (this is the aperture problem illustrated in Fig. 11). To recover the full velocity field requires constraining the possible range of solutions. Hildreth proposes a smoothness constraint, which is based on the physical assumption that surfaces are generally smooth. A theoretical analysis of the conditions under which these assumptions yields the correct velocity field has been completed (Hildreth, 1984). An algorithm has also been devised and implemented, using several examples that permit comparison with psychological observations. Examples of particular interest are the rotating spiral and the barber pole illusion, where the true velocity vectors are not seen either by the algorithm nor by the human observer (Fig. 12). Such failures consistent with human perception give credibility to Hildreth's formulation, suggesting it may be the basis for motion analysis in the most powerful vision machine known—our own!

6.0 Summary

All plane curves may be described by a linked list of codons. For smooth, closed plane curves, the codon sequences enable us to enumerate shapes of increasing complexity. The "ellipse" "peanut" and "dumbbell" are the simplest shapes according to the theory. Even these simple shapes, when embedded in a hierarchy, provide very powerful indices

for object recognition.

A nice feature of the codon theory is that the descriptions are computable from images, as we have shown. Current work in progress is the creation of a completely automatic package that will deliver a hierarchy of codon descriptions for a 512 x 512 static image. At a later date we will add a motion capability for single blob isolation along the lines Hildreth has proposed.

At present, theoretical work has been concentrating on what assertions about 3D shape can be made from the 2D codons. Metric information required for a more detailed description of a "part" is also being considered. Possibilities include abstracting qualitative descriptors of part-shapes such as "knob", "neck", "bump", "dent", "fold", "finger", etc. The relations between blobs and their sub-blobs must also be made explicit. However, at present we have little insight how this should be done, with the exception of Ullman's paper on "Visual Routines", together with some crude psychophysical observations that we have tentatively begun to explore. Over the next year, therefore, our primary effort will be toward extending the codon theory, and its implementation.

7.0 Appendix 1: Software Developed

Over 200 programs have been written so far for the work of this grant. The programs range from simple subroutines to major packages and systems programs. Here we list the major classes of programs and describe a few examples.

PACKAGES

Pyramid: This code generates the Gaussian Pyramid from an input image. We have three versions of this code. The original version was developed on a Symbolics 3600 Lisp machine, the next version on our VAX, and the final version is in Adage 3000 BPS microcode. Since this routine is basic to our work, we feel that the speed improvement (a factor of 10) in the final version was worth the effort.

Blob: A package for generating the first level symbolic blob descriptors from a Gaussian Pyramid. The images are made into binary images, and these are used to generate a tree structure of blobs. A blob is represented by its location, size, outline, tangents, and pyramid level. Additional constraints, for example color and motion, on the definition of the blobs will be used to reduce the number of blobs and improve their definition.

Codon: This package takes the edge list generated by the BLOB package and generates a Codon description. It differentiates the curve tangents, and applies algorithms to define extrema of curvature. Additional programs are used for display, test, and rudimentary matching of symbolic descriptions.

Cartoon: We originally wrote the CARTOON package on the PDP 11/23. This re-write of the package takes advantage of the VAX's increased performance. The CARTOON package has a user-friendly interface and allows users to manipulate images by doing convolutions, thresholds, color coding, acquiring, saving, and restoring, etc.

Fractals: This package allows the generation of a limited class of fractal patterns. We are interested in these patterns as potential descriptors of form and texture in complex images.

Ohio: We have modified the Ohio Rendering Package provided by Prof. Frank Crow to run on the Adage 3000. This package allows us to generate realistic objects with shaded surfaces, multiple light sources, etc. Used for studying the perception of objects.

Spheres: A package for generating spheres with various shading, reflectance, and light-source functions. Used for studying the perception of reflectance and specularly.

SYSTEMS PROGRAMS

We have written drivers, linkers and controllers to effectively use the new hardware acquired in part on this grant.

IK: This is the Adage 3000 (Ikonas) system driver. This is a modification of code originally provided by Ron Gordon of Bell Laboratories. Our code is now a distributed Berkeley 4.2 Unix driver for this machine.

RILMOD: We use the GIA/RIL microcode assembler for the Adage 3000 (provided by the University of North Carolina). Extensive modifications of the RIL linker have been made to improve performance and support our hardware.

MC68: Code to utilize the Motorola 68000 processor on the Adage 3000. This enables us to load and run programs, handle interrupts and access the other hardware on the Adage from the 68000.

SLAVE MACHINES

Our PDP-11's may be used either as stand-alone machines for image collection and processing, or as "slave" machines performing specific functions under the direction of the VAX.

SEND11, GRAB11, IMAGE11: Allow the VAX to control a PDP-11 used as a remote graphics and image processor.

SLAVE: A remote graphics and image processing that runs on the PDP-11's. Performs such functions as drawing lines, text, circles, etc., acquiring and displaying images.

LIBRARIES

A variety of libraries have been written to support the research program. A library is a collection of routines that serve a common goal.

LIKNEW: The New Ikonas library contains routines for controlling the Adage 3000, graphics (lines, circles, blocks, text, etc.), image acquisition, save and restore, etc. This library contains over 60 subroutines, and includes an extensive shared data base for specification of the Adage parameters.

GRAPHICS: A collection of simple and advanced graphics routines written in portable C code so they can be used on a variety of machines.

USEFUL: A collection of useful tools and routines. Used in most of our programs.

COMMUNICATIONS

Our VAX is linked (currently via RS-232 lines) to the main AI machine (OZ) and our "slave" PDP-11's. These connections will be replaced by Ethernet links.

URT, RTU and UNET: Communication and file transfer between the VAX and a PDP-11. Allows a PDP-11 to appear as a "virtual" VAX terminal. This is an extensive re-write of code provided by the Center for Cognitive Science at M.I.T.

PR: Remote printing of text files on various printers.

OZLINK, OZTALK: Communication and file transfer (using the KERMIT protocol) between the VAX and the OZ (AI) machine.

TESTS AND DEMONSTRATIONS

TEST programs include programs for testing the Adage 3000's various components, for testing the communication links, and for testing developing code. The DEMONSTRATION packages link components and other packages to demonstrate such things as the entire blob/codon system and the graphics capability of the Adage 3000.

Appendix 2: Publications and References

- Bennet, J.B. & Hoffman, D. "Shape decomposition via transversality". *Image Understanding* (1985).
- *Bobick, A. "Grouping visual targets". Abstract. *Perceptual Organizations Workshop, Pajaro Dunes, CA, 1984.*
- *Bobick, A. "Using mirror reflections to recover shape". Abstract. WH2, OSA Annual Meeting, 1984.
- Burt, J.B. & Adelson, E.H. "The Laplacian Pyramid as a compact image code". *IEEE Trans. on Comm.*, Vol. COM-31, No. 4, pp. 532-540, April 1983.
- Burt, J.B. "The pyramid as a structure for efficient computation". *Rensselaer Polytechnic Institute Image Processing Laboratory Technical Report IPL-TR-038, 1982.*
- *Dawson, B.M. & Treese, G. "Computing curvature from images". *Proceedings SPIE*, vol. 504, November, 1984.
- *Dawson, B.M. & Treese, G. "On finding objects in images". *SPIE Conference on Image Processing*. Accepted for publication, January, 1985.
- *Gilson, W. "Idiot's guide to OZ (A manual for text editing)". *Natural Computation Group*, 48 pages (1984).
- *Hildreth, E.C. "The computation of the velocity field". *MIT AI Memo 734, 1983.*
- *Hildreth, E.C. "Computations underlying the measurement of visual motion". *MIT AI Memo 761, (1984).*
- *Hildreth, E.C. "Computing the velocity field". Abstract. TuC2, OSA Annual Meeting, 1984.
- *Hildreth, E.C. *The Measurement of Visual Motion*. MIT Press, Cambridge, MA, 1984.
- Hoffman, D. & Richards, W. "Representing smooth plane curves for visual recognition: implications for figure-ground reversal. *Proc. American Assoc. Artificial Intelligence (AAAI-82)* (1982).
- *Hoffman, D.D. & Richards, W. "Parts of recognition". *MIT AI Memo 732 (1983).*
- *Hoffman, D.D. & Richards, W. "Parts of recognition". Abstract. TuC3, OSA Annual Meeting, 1984.
- *Hoffman, D.D. & Richards, W.A. "Parts of recognition". *Cognition*, in press.
- *Supported by current AFOSR contract.

- Koenderink, J.J. "The structure of images". *Biol. Cybern.*, 50:363-370 (1984).
- Richards, W., Nishihara, H.K. & Dawson, B. "Cartoon: a biologically motivated edge detection algorithm". *MIT AI Memo 668* (1982).
- *Richards, W. "Structure from stereo and motion". *MIT AI Memo 731* (1983).
- *Richards, W. "Interpreting image contours". Abstract. Air Force Vision Program Review, May 1984.
- *Richards, W. & Hoffman, D. "Codon Constraints on Closed 2D Shapes". *MIT AI Memo 769* (1984).
- *Richards, W. "Grouping without prior correspondence". Abstract. WP1, OSA Annual Meeting, 1984.
- *Richards, W. "Structure from stereo and motion". *J. Opt. Soc. Amer.*, January 1985. (Also *MIT AI Memo 731* (1983)).
- Rubin, J. & Richards, W. "Color vision and image intensities: when are changes material?" *Biol. Cybern.*, 45:215-226 (1982).
- *Rubin, J. & Richards, W. "Color vision: representing material changes". *MIT AI Memo 764* (1984).
- *Rubin, J. "Telling actors from objects". Abstract. WP4, OSA Annual Meeting, 1984.
- *Ullman, S. "Visual Routines". *Cognition*, in press (1985). [Also *Center for Cognitive Science Occasional Paper 20* (1984).]
- Yuille, A. & Poggio, T. "Fingerprint theorems for zero-crossings". *MIT AI Memo 730* (1983).
- *Yuille, A. & Poggio, T. "Fingerprint theorems". *Proc. Assoc. Artificial Intelligence (AAAI-84)* (1984).

END

FILMED

4-85

DTIC