UTEC-CSc-77-118
Final Technical Report
Computer Science
April 1977

12

SENSORY INFORMATION PROCESSING

UNIVERSITY OF UTAH

Sponsored by
Defense Advanced Research Projects Agency
ARPA Order Number 2477

Approved for public release;
distribution unlimited.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency of the U. S. Government.

SENSORY INFORMATION PROCESSING

1 July 1976 - 31 March 1977

Final Technical Report

Contractor:	University of Utah
Contract Number:	DAHC15-73-C-0363
Effective Date:	1 July 1973
Expiration Date:	31 March 1977
Amount of Contract:	$2,825,000.00
Project Code:	3D30

Principal Investigator:	Dr. Thomas G. Stockham, Jr.
Telephone:	(801) 581-8224
Contracting Officer:	Mr. Edgar S. Allen
	DSSW

Approved for public release;
distribution unlimited.

Sponsored by
Defense Advanced Research Projects Agency
ARPA Order Number 2477

UTEC-CSc-77-118

(See 1473)

TABLE OF CONTENTS

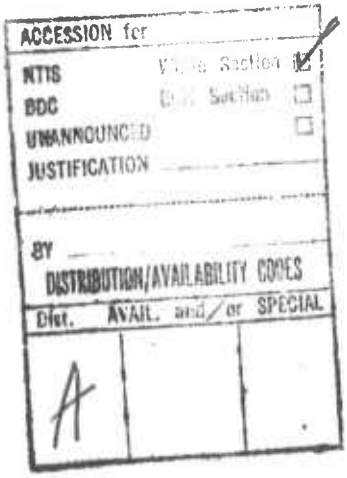

REPORT SUMMARY

This final report covers the nine month period from July 1976 through the termination of the contract at the end of March 1977. As such, this report describes the accomplishments of this period in detail and thus includes the semi-annual report of the period July 1976 - December 1976.

Section I discusses the theoretical underpinnings of a new mathematical model describing some of the perceptual characteristics of human sensory information processing. While much of the mathematics of the proposed model is explained here, experimental testing and verification has yet to take place. This work will continue to be pursued under alternate funding.

Section II discusses in detail the background considerations necessary to understand the problems of removing atmospheric turbulence blur from images. This area of research is currently in transition from the stage of artificial computer models, to the stage of working with the simplest real data, i.e. case of a point light source, a star, viewed through the earth's atmosphere. After the termination of this contract this work will be funded by the Air Force Office of Scientific Research under grant number 77-3212.

Section III reports the background and progress to date on building the tools necessary for an image understanding system. These tools as described were incorporated into a manual image understanding system designed to do analysis by synthesis. The system handles two-dimensional single objects on non-textured backgrounds, but only at the level of a trial testbed. Necessary next steps in developing a working system are outlined.

SECTION I

GROUP REPRESENTATIONS AND THE MODELING OF VISUAL PERCEPTION

James T. Kajiya

The student of human visual perception is often overwhelmed by the vast amount of data that has been accumulated from experiments performed within the last century or so. It is often difficult to understand why a certain experiment has been performed. Results from similar experiments sometimes seem to conflict. Further confusion results when the student encounters raging controversies, the resolution of which would seem to minimally advance our knowledge of how we see. The reason for all this trouble stems from the fact that a suitable superstructure providing organization and support of this accumulation of data does not exist, i.e. an adequate theory of perception is not in hand.

Theories, by their nature, are primarily qualitative descriptions of what is going on. Of the qualitative observations that can be made about the human visual system, some of the most fruitful are those of the perceptual constancies. In short, psychophysicists have puzzled over the eye's marvelous capacity to respond gracefully over wide

transformations on the images presented to the retina. The perception of a triangle is relatively unaffected by variations in size, orientation, and perspective. This particular phenomena takes as its name "shape constancy". These constancies or invariances provide enough qualitative information to draw upon the powerful theory of group representations.

This paper will report on some first steps taken toward the application of the theory of group representations to the modeling of human perception. In particular, Section 2 will briefly discuss two groups that express certain observed invariances within the visual system. Section 3 will perform the relatively straightforward calculations of the unitary equivalence classes of irreducible representations for one of these groups. Section 4 will attempt to carry out a brute force harmonic analysis. It will turn out that this brute force approach will fail in that it will not shed any light upon the workings of the visual system. This may explain in part why Abstract Harmonic Analysis has not seen wide use in the study of vision to date. Section 5 (the principal contribution of the paper) will examine why the brute force approach fails. Also an operator imbedding the image space into a larger space is presented for the first time. This new operator will be seen to possess certain attractive properties that may turn out to make it quite useful as a gadget for image

processing in general. Finally, Section 6 will quickly outline the rest of the program needed to obtain a perceptually important transform that will mimic the operation of the visual system.

2. THE PERCEPTUAL GROUPS

It has been known for quite some time that the groups describing the various perceptual constancies are Lie Groups [1, 2]. In particular, Pitts and McCulloch in [1] conducted a kind of harmonic analysis at the "DC" frequency. The main group that should be kept in mind is the group G of affine transformations on $\mathbb{R}^2$. Our group G expresses the form constancy in the literature. Let (x,y,w) be homogeneous co-ordinates of the plane. The one parameter subgroups generating G are:

translations $\begin{pmatrix} 1 & 0 & a \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix}$

rotations $\begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}$

dilatations $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^r \end{pmatrix}$

Notice that if the identification of $\mathbb{R}^2$ with the complex plane is made then G is the complex "$\alpha x+\beta$" group where

$$\beta = (a,b)$$

$$\alpha = e^{r+i\theta}$$

Since the first derived subalgebra $[\mathcal{G},\mathcal{G}]$ of the Lie Algebra $\mathcal{G}$ of G is the set of translations which is abelian, this group is solvable. It is also apparent from the $\alpha x+\beta$ expression that G is the semi-direct product of the two abelian groups $(\mathbb{C},+)$ and $(\mathbb{C}\setminus\{0\},\cdot)$.

There is some reason to suspect that rotations do not correspond to a perceptual invariance [3, 4]. Furthermore bilateral symmetry seems to evoke a strong mechanism in the visual system. Indeed, study of this kind of symmetry dates back to Mach. Thus another important group we should consider is generated by translation, dilatation, and

$$\text{reflection} \qquad \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

This group is also a semi-direct product of two abelian groups. In what follows we treat only the first group.

Since these groups have such a nice structure, it is quite straightforward to apply the Mackey Induction Theory to obtain a complete characterization of their unitary equivalence classes of irreducible representations $\hat{G}$.

3. CALCULATION OF THE DUAL $\hat{G}$

Let us recall some basic definitions and theorems [5, 6]. Definition. Let γ be a representation of a subgroup H of G on a hilbert space $\mathcal{H}_\gamma$. Consider the space $\mathcal{H}(\gamma)$ of all functions from G to $\mathcal{H}_\gamma$ satisfying

$$\text{i) } x \to (f(x),\xi) \text{ is Borel } \forall \xi \in \mathcal{H}_\gamma$$

$$\text{ii) } f(hx) = \gamma(h)f(x),\ h\in H, \text{ almost all } x\in G$$

$$\text{iii) } \int_{G/H} ||f(x)||^2 \, d\bar{x} < \infty$$

where we make the usual L_2 identification. Set

$$\pi(g)f(x) = f(xg)\,(\rho)^{1/2} \quad f\in\mathcal{H}_\gamma$$

and call π the representation of G induced from H by γ. We write $\pi = \mathrm{Ind}_H^G(\gamma)$.

Theorem. Let G be a regular semi-direct product of H and N, with N abelian and normal. For each orbit $H\cdot\gamma$, $\gamma\in\hat{N}$, set $H_\gamma=\{h\in H: h\cdot\gamma=\gamma\}$ Let $U_{(x,y)}=V_xW_y\in\hat{G}$. Then the projection-valued measure defined by W in $\hat{N}$ is concentrated

in a single orbit $H\cdot\gamma$ and $V = \mathrm{Ind}_{H_\gamma}^{H}(\nu)$ for some $\nu\varepsilon H_\gamma$. Every pair consisting of an orbit $H\cdot\gamma$ and irreducible ν arise in this way, and two of these are equivalent if the orbits are identical and the representations ν are equivalent.

The above definition and theory were taken from the excellent [7].

Since both perceptual groups have available the abelian normal subgroup $\mathbb{R}^2$, the characters in question γ are the familiar fourier kernels which are isomorphic to $\mathbb{R}^2$. The orbits form two disjoint subsets of the plane: The point (0,0) and the punctured plane. Their respective isotropy subgroups are H itself and the identity.

4. A NAIVE ATTEMPT

Since we are interested in images, the hilbert space that our representation acts upon will be $\mathcal{H} = \mathbb{L}_2(\mathbb{R}^2, \mu)$. The action of G on $\mathcal{H}$ will be

$$U_g(f(x)) = \alpha(g)f(gx)$$

The multiplier $\alpha(g)$ is the square root of the Radon-Nikodym derivative of the measure $\mu_g(E) = \mu(gE)$ with respect to

Lesbegue measure μ. It is necessary because, as a moments reflection will show, the best we can do is quasi-invariance of our measure. The α of course makes U a unitary representation.

Our aim is to decompose this representation into irreducible subrepresentations and then reap the benefits by identifying the projection-valued measures. Unfortunately this representation is itself irreducible so our hopes of arriving at a decomposition have evaporated. The irreducibility of this representation indicates that any ad-hoc scheme that attempts to "crank in" rotations and dilatations will meet with only partial success at best. The main problem with our space of images is that there is not enough "elbow room" to be able to proceed with our decomposition. To see this we present a "heuristic" argument. Let f(x) be the characteristic function of, say, the unit square. We can approximate any characteristic function of a measurable subset of the plane simply by following the classical Lesbegue construction. The linear closure of these functions form the simple functions. Since the simple functions are dense in $\mathbb{L}_2$ we see that any picture is related to any other.

5. THE MANDALA TRANSFORM

Obviously something must be wrong with our model of the human visual system. The eye can successfully decompose images into separate parts. The discrepancy can be accounted for by the notion of texture. Key to the above heuristic argument was the freedom to shrink the unit square to arbitrarily small proportions. If this is done with an image then at some point the image will change its perceived character. Objects in the image will become so small that they will become indiscernible and turn into texture. A beautiful example of this effect are the images by Harmon and Knowlton [8]. In other words the dilatation operation is not a true symmetry operation. How then can we have form constancy?

Further clues to the operation of the visual system can be culled from the landmark experiments of Hubel and Wiesel [9]. In these experiments the single neurons in the visual cortex monitored by microelectrodes were found to respond only to small portions of the visual field. An examination of several closely spaced neurons indicated slightly displaced but overlapping "windows" into the visual field. Furthermore, many neurons were responsive only to lines placed in a certain orientation with respect to the window.

Hubel and Wiesel concluded that they had found mechanisms closely resembling "line detectors" within the visual cortex. Further experiments by others [10, 11] indicated that these neurons were not line detectors but rather spatial frequency filters. Some models have already used this fact [12].

The above discussion motivates the following:

Definition. The Mandala transform is an operator $M: L_2(\mathbb{R}^2,\mu) \to L_2(\mathbb{R}^4,\mu)$. Let h be a smooth function such that $h(0) = 1$ then

$$Mf(x_1,x_2,\omega_1,\omega_2) = \hat{f}(x,\omega) =$$

$$\int_{\mathbb{R}^2} h(\xi)f(\xi-x)\exp(-i\langle\xi,\omega\rangle)d\mu(\xi)$$

where $x,\omega \in \mathbb{R}^2$. This is a quite familiar object in the engineering circles around speech processing [13, 14] where, of course, functions on the line instead of on the plane are considered. The weighting function h plays the role of our window. Indeed, in the engineering literature, h is called the window function and is usually chosen to be a Hamming window or Hanning window. If x is fixed then we can see that the operation of the transform is to first position f under the window and then fourier transform the product.

The inversion formula is

$$f(x) = (2\pi)^{-2} \int_{R^2} \hat{f}(x,\omega) \exp(i\langle x,\omega\rangle) d\mu(\omega)$$

To define a representation in the new space we first need some simple facts.

Proposition. Let T be a non-singular linear transformation of the plane. Then

$$(f_T)\hat{}(x,\omega) = f(Tx,(T^*)^{-1}\omega) \frac{1}{\det T}$$

where $\hat{f} = Mf$ and $f_T(x) = f(Tx)$.

Proof. This follows by substituting into the inversion formula and using the adjoint at the inner product.

Proposition. $(f_\xi)\hat{}(x,\omega) = \hat{f}(x+\xi,\omega)e^{i\langle \xi,\omega\rangle}$

where $f_\xi(x) = f(x+\xi)$.

Proof. Similar.

The new representation is now defined accordingly. Notice that for a rotation 0 we have $(0^*)^{-1} = 0$ since 0 is real unitary.

Why will our new representation be any different than the old? This question points out a number of interesting effects that arise from the Heisenberg uncertainty principle. In particular, the uncertainty principle has a lot to say about the "shape" of points in our new space $L_2(\mathbb{R}^4,\mu)$ that are images under the operator M [15]. Thus, because some functions violate this shape criterion the old space is embedded within the new space and can be decomposed in a way not possible before.

The uncertainty principle also elucidates mathematically the intuitive notion of texture. In our new space an image is represented as a function $\hat{f}(x,\omega)$ where both $x, \omega \in \mathbb{R}^2$. In the following sense x is a "feature" or "object" space while is a "texture" space.

Features in a given image can become texture when seen from afar. Similarly when one examines a texture pattern that is suitably magnified, features appear. These of course are very imprecise notions. Furthermore, the classification of an image into texture or feature is quite easy for extremes but questionable otherwise. Consider a texture pattern and begin zooming in on it. The perceived change is gradual. Any abruptness really stems from the language we use to describe these things. The uncertainty

principle under this zooming experiment will cause an image to slowly rotate out of the feature space and into the texture space.

Another attribute of the Mandala transform space is that it represents both local and global aspects of the image. (These terms are meant in the engineering sense.) That a successful visual model must necessarily represent both aspects of an image and of the complex interactions between them can be seen by citing the Cornsweet illusion, or the phenomenon of subjective contours [16, 17] in which strictly local or texture aspects will cause striking global illusions.

It must be mentioned also that the Mandala transform bears a distinct resemblance to several procedures well known to engineers in the Image sciences [18, 19]. It is somewhat curious this transform is not in common use, given its yeoman service to the speech community.

6. HARMONIC ANALYSIS

Now that the representation is defined and the dual $\hat{G}$ is in hand we have only to construct a direct integral decomposition [20]. Instead of a factor decomposition we

really need Mautner's adaptation [21]. However, since our groups are of type Lie we have a method for finding the maximal subalgebra of the commutant. We follow Bargmann [22] by finding the center of the associated Universal Enveloping algebra and constructing the derived representation with the aid of the theorems of Gårding, Segal, Nelson and Stinespring [23, 24, 25, 26]. Thus we are able to construct the "equations of vision" similar to the field equations of particle physics [27]. Solving these generalized Laplace equations we obtain the kernels for the desired projection operators.

In this way we are able to construct a perceptually important transform such that a modification in the transform space becomes a perceptual correlate.

REFERENCES

[1] W. Pitts, W.S. McCulloch, "How We Know Universals: The Perception of Auditory and Visual Forms", Bull. Math. Biophysics 9 (1947) pp. 127-147.

[2] W.C. Hoffman, "The Lie Algebra of Visual Perception", J. Math. Psych. 3(1966) pp. 65-98.

[3] I. Rock, "The Perception of Disoriented Figures" in Readings from Scientific American: Image, Object, and Illusion.

[4] R. Shepard, J. Metzler, "Mental Rotation of Three-dimensional Objects", Science 171 (1973) pp. 701-703.

[5] G.W. Mackey, "Induced Representations of Locally Compact Groups I", Ann. of Math. 55(1952) pp. 149-204.

[6] ----------, "Induced Representations of Locally Compact Groups II: The Frobenius Reciprocity Theorem", Ann. of Math. 58(1953) pp. 193-221.

[7] R.L. Lipsman, Group Representa ions, Lecture Notes in Math. No. 338, Springer Verlag (1974).

[8] L.D. Harmon, K. Knowlton, "Picture Processing by Computer", Science 164(Ap. 4, 1969) pp. 19-29.

[9] D.N. Hubel, T.N. Wiesel, "Receptive Fields and Functional Architecture of the Monkey Striate Cortex", J. of Physiology (London) 195(1968) pp. 215-244.

[10] F.W. Campbell, J.G. Robson, "Application of Fourier Analysis to the Visibility of Gratings", J. of Physiology (London) 197(1968) p. 551.

[11] F.W. Campbell, F.F. Cooper, C. Enroth-Cudgell, "The Spatial Cells of the Cat", J. of Physiology (London) 203(1969) p. 223.

[12] B. Baxter, "Image Processing in the Human Visual System", Computer Science Tech. Report. UTEC-CSc-75-168, Univ. of Utah, December 1975.

[13] M. Callahan, "Acoustic Signal Processing Based on the Short Time Spectrum", Computer Science Tech. Report, UTEC-CSc-76-209, Univ. of Utah, March 1976.

[14] M. Portnoff, "Implementation of the Digital Phase Vocoder Using the Fast Fourier Transform", MIT preprint 1975.

[15] D. Slepian, H. Pollak, "Prolate Spheroidal Wave Functions, Fourier Analysis and Uncertainty II", Bell Sys. Tech. J. 40(1961) pp. 65-85.

[16] T. Cornsweet, Visual Perception, Academic Press, 1970.

[17] G. Kanizsa, "Subjective Contours", Scientific American, April 1976.

[18] T.G. Stockham, Jr., T.M. Cannon, R. Ingebretsen, "Blind Deconvolution Through Digital Signal Processing". Proc. IEEE, vol. 63, no. 4 (1975).

[19] G. Anderson, T. Huang, "Piecewise Fourier Transforms for Picture Coding", IEEE Trans. Comm. Tech. COM-9(1971), pp. 133-140.

[20] J.V. Neumann, "On Rings of Operators. Reduction Theory", Ann. of Math. 50(1939), pp. 142-204.

[21] F. Mautner, "Unitary Representations of Locally Compact Groups", Ann. of Math. 51(1950) pp. 1-25.

[22] V. Bargmann, Irreducible Representations of Lorentz Group", Ann. of Math. 48(1974) pp. 568-640.

[23] L. Gårding, "Note on Continuous Representations of Lie Groups", Proc. Nat. Acad. Sci. USA, v. 48(1947).

[24] I.E. Segal, "Hypermaximality of Certain Operators on Lie Groups", Proc. Am. Math. Soc. v. 9(1952) pp. 13-15.

[25] E. Nelson, W. Stinespring, "Representation of Elliptic Operators in an Enveloping Algebra", Am. J. Math. 81(1959), pp. 547-560.

[26] E. Nelson, "Analytic Vectors", Ann. of Math. 70(1959) pp. 572-615.

[27] V. Bargmann, E. Wigner, "Group Theoretical Discussion of Relativistic wave equations", Proc. Nat. Acad. Sci. USA 34 (1948) pp. 211-223.

SECTION II

REMOVING ATMOSPHERIC TURBULENCE INDUCED BLUR FROM ASTRONOMICAL IMAGES

Brent Baxter

Craig Rushforth

1. Introduction

One of the most difficult and serious problems associated with high-resolution optical imaging systems is the distortion introduced by atmospheric turbulence. The bulk of this distortion is due to random phase perturbations caused by variations of the refractive index along the propagation path. If images produced in the presence of atmospheric turbulence are simply recorded and processed by conventional methods, the result is often a loss of essential high-spatial-frequency detail. This problem is serious even under favorable conditions at a good site, and it is made worse by such effects as wind, solar heating, and boundary-layer turbulence near a moving object. Much of the potential resolving power of a large telescope will be lost unless the received optical field is recorded and processed properly.

The removal of the effects of atmospheric turbulence from optical images is a problem of long standing. Recently three new techniques have been developed which could potentially lead to significant improvements in the resoultion achievable when imaging through a turbulent atmosphere. In Section 2, we provide some historical background and then discuss these recent developments in some detail. Section 3 contains a description and discussion of some one-dimensional simulations which we

performed to test the theory and to provide an indication of the effects of sensor noise on the performances of these procedures. In Section 4, we summarize our results, discuss the current status of the problem, and provide some guidelines for further research.

2. Summary of Previous Work

In this section, we provide a brief summary of previous work on the problem of restoring images which have been blurred by atmospheric turbulence. We make no attempt to be comprehensive, but rather restrict our attention to work which is immediately relevant to our problem. For a more complete discussion of other work, and for a comprehensive bibliography, see [1].

2.1. Long-Time Averaging

The problem in which we are interested is characterized by an object intensity distribution which is constant for a long period of time, during which the conditions of the atmosphere change appreciably. One obvious approach to removing the effects of random atmospheric turbulence is to calculate a long-time average of the image intensity distribution. This type of averaging can be effected either by exposing film or another recording medium for a long time, or by averaging many short-time exposures. In either

case, averaging has the effect of removing the random fluctuations from the resulting image. Figure 1 shows such an image.

Unfortunately, this long-time averaging also has the effect of removing or greatly attenuating the high spatial frequencies from the image [2]. Thus, although the average image is stable, its spatial-frequency content is limited by atmospheric "seeing" rather than by the imaging system. Furthermore, only limited improvement can be achieved by applying standard restoration techniques to the long-time average image [3]. As an example of the seriousness of this effect, consider the 200-inch Hale telescope. The diameter of the Airy disk for this telescope is about 0.05 arc seconds, while a typical seeing limit on a quiet night is about 2 arc seconds. Thus, if long-time averaging is employed, atmospheric turbulence degrades the resolving power of this system by a factor of about 40. If high-resolution imaging systems are to reach their potential in the presence of atmospheric turbulence, then, it is clear that alternative methods are required which yield statistical stability while preserving high spatial frequencies.

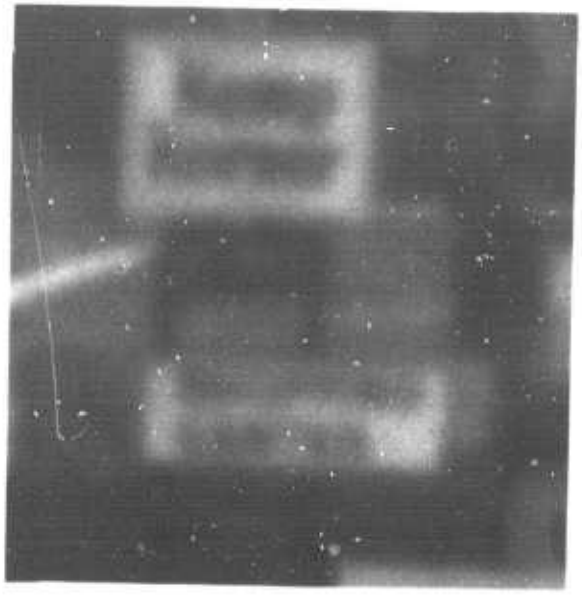

Figure 1. Image blurred by atmospheric turbulence. Two second time exposure made through a questar telescope.

2.2. Interferometric Methods

Interferometers have been used in high-resolution astronomy since the beginning of the century. In 1920, Michelson [4] and Anderson [5] showed that a stellar interferometer could be used to resolve objects which were normally beyond the seeing limit of the atmosphere. Although the stellar interferometer is theoretically capable of measuring both the amplitude and the phase of the coherence function and thereby completely recovering the object intensity distribution, in practice the phase cannot be accurately determined. This limits the usefulness of the stellar interferometer to regular objects such as binary stars whose Fourier transforms have well-defined zeroes. Furthermore, the stellar interferometer is inherently a one-dimensional scanning device whose application to two-dimensional objects would be cumbersome and expensive.

A modified interferometer known as the intensity interferometer was introduced by Hanbury-Brown and Twiss [6] to improve upon the performance of the stellar interferometer in certain situations. This device is inherently restricted to measuring the amplitude of the coherence function, and like the stellar interferometer is basically a one-dimensional instrument. Thus, for our purposes, it suffers from essentially the same limitations as does the stellar interferometer. Hence, alternative

procedures are needed for those cases in which the object intensity distribution is asymmetric and in which complete two-dimensional information about the object is desired.

A number of modifications to the basic stellar and intensity interferometers have been suggested over the years. We mention two of these proposed modifications very briefly. For a more comprehensive discussion, see [1].

Since the intensity interferometer by itself cannot measure phase, it must be modified or augmented if the complete object is to be reconstructed. If a known object such as a point source is present in the field of view along with the unknown object, the phase of the unknown object can be obtained from the cross term resulting from interference with the reference object. This is entirely analogous to the use of a point source to retrieve phase in a hologram. Although this is an interesting idea, its application to real astronomical problems is obviously limited because of the required presence of a known reference object.

An interesting technique which combines features of the stellar interferometer with those of the intensity interferometer has been suggested by MacPhie [7, 8]. In principle, this device can measure both amplitude and phase, and can therefore reconstruct the object intensity. Despite its apparent advantages, however, this technique has not been applied to real data. Furthermore, like its

predecessors the stellar interferometer and the intensity interferometer, it is inherently a one-dimensional system whose adaptation to extended two-dimensional objects would be cumbersome at best.

2.3. Speckle Interferometry

A partial solution to the problem of reconstructing two-dimensional objects was described by Labeyrie [9, 10], whose technique has come to be known as speckle interferometry because of the speckled character of the images used. An example of such an image is shown in Figure 2. In this technique, a series of short-time exposures of the atmospherically-degraded object are taken through a narrow-band spectral filter. The exposure time is short enough that the atmosphere is effectively "frozen" during the exposure, and successive photographs are sufficiently separated in time that they may be taken to be statistically independent for the purpose of averaging.

Instead of simply averaging these short-time exposures directly, a process which leads to a loss of high spatial frequencies as we have seen, speckle interferometry requires some additional operations before averaging. First, the Fourier transform of each image is taken, either optically or digitally. The averaging is then performed on the squared magnitude of these individual Fourier transforms.

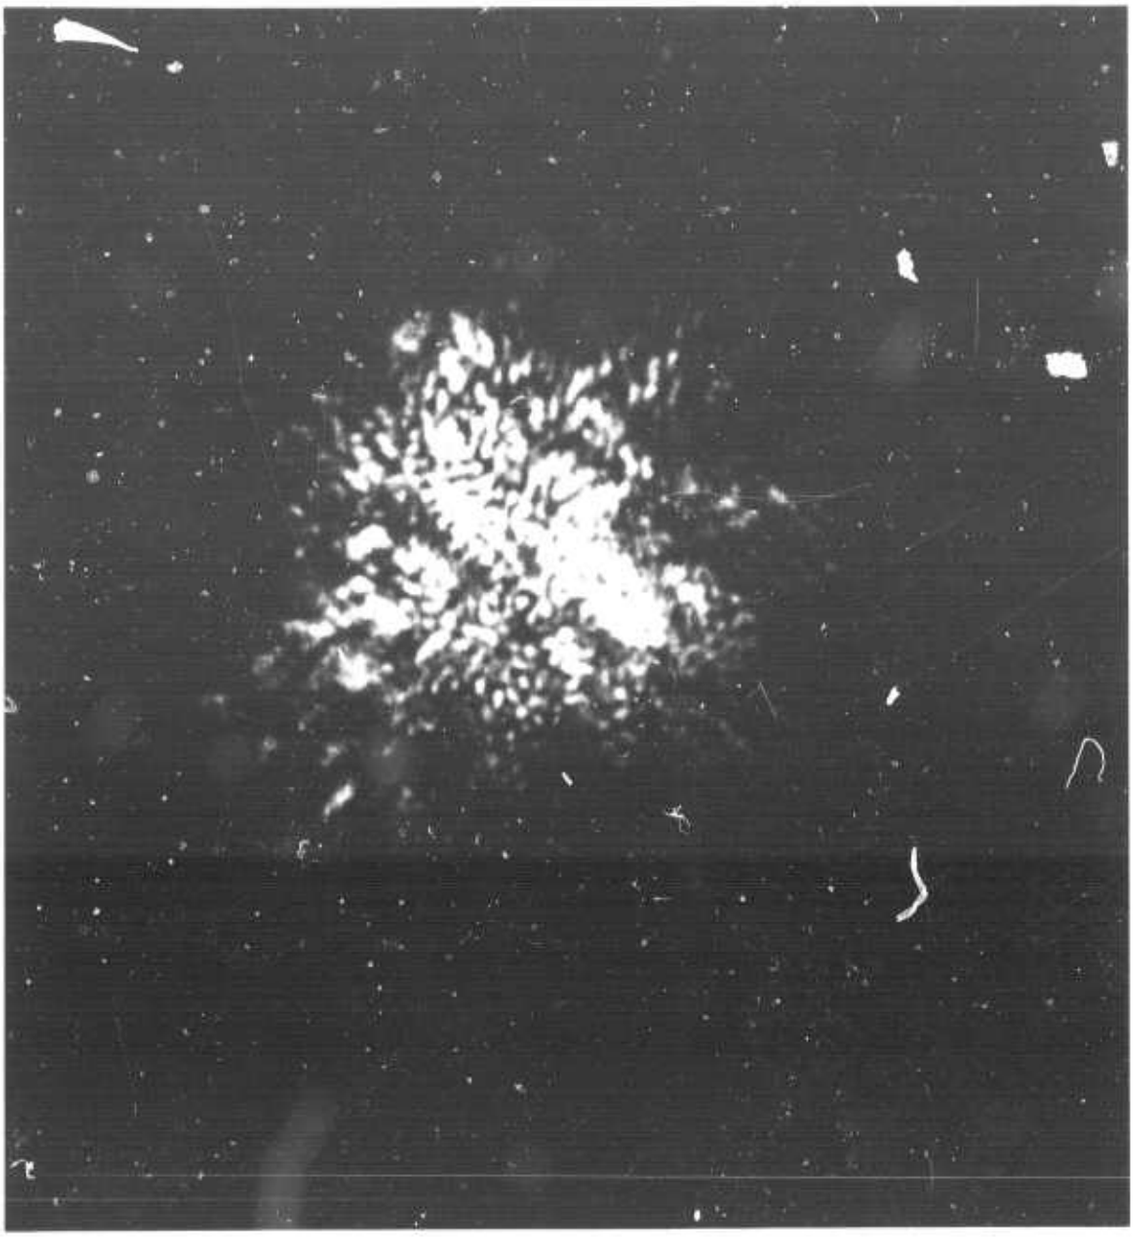

Figure 2. Speckle image of Arcturus made on the Mayall four-meter telescope at Kitt Peak by S. D. Worden and B. Baxter Delays imposed on the wavefront by its passage through the atmosphere cause speckles whose size is characteristic of the diffraction limit of the telescope.

The process of taking the square of the magnitude destroys the information regarding the phase of the image. In this fact lies both the strength and the weakness of speckle interferometry. Since it is the random phases of the high-spatial-frequency components which cause their mutual cancellation in the process of long-time averaging, the suppression of phase prevents this cancellation and thus preserves some high-frequency information which would be lost under direct long-time averaging. Unfortunately, information about the phase of the object is also suppressed. This loss of phase information limits the class of objects to which speckle interferometry can be usefully applied. For example, this technique can yield useful information about symmetric objects or binary star systems, but not about an object with an arbitrary intensity distribution or arbitrary shape.

Labeyrie and his colleagues have verified experimentally that speckle interferometry yields essentially diffraction-limited information about such objects as binary star systems. In addition, theoretical discussions of this technique have been provided by Miller and Korff [11] and by Korff [12]. These authors have used models of propagation in a turbulent atmosphere to explain why speckle interferometry works and to generalize the results to partially-coherent objects.

2.4. The Knox-Thompson Technique

The success of speckle interferometry in reducing the effects of atmospheric turbulence, coupled with its limited applicability, has spurred a search for generalizations and improvements of this technique. A significant extension of speckle interferometry was developed by Knox and Thompson in 1974 [1]. This new procedure uses speckle interferometry as before to obtain amplitude information. To this amplitude information is added phase information obtained by averaging and processing the Fourier transforms of the short-time exposures in a different way.

In this section we present a brief discussion of the Knox-Thompson technique. This is not intended to be a complete treatment, but merely an attempt to cover the essential features of the procedure in order to make the present report reasonably self-contained. For a more thorough treatment, see [1].

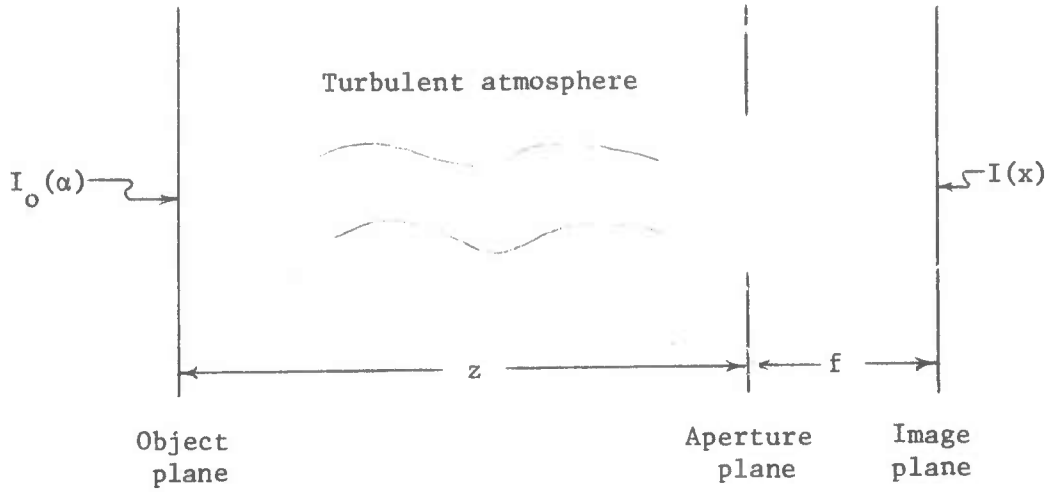

Figure 3. Imaging configuration

The imaging geometry is shown in Figure 3. Before proceeding with a discussion of the Knox-Thompson technique, we define the following quantities:

$I_o(\alpha)$:	Object intensity distribution.
$V(\hat{x})$:	Complex wavefront in the aperture plane due to point source at origin in object plane.
$\phi(\hat{x})$:	Phase of $V(\hat{x})$.
$A(\hat{x})$:	Pupil function of imaging system.
$I(x)$:	Image intensity distribution.

$S(x)$: Instantaneous (short-exposure) point spread function of atmosphere-telescope combination.

$\tilde{I}_o(u)$, $\tilde{I}(u)$, $\tilde{S}(u)$: Fourier transforms of I_o, I, and S, respectively.

f: Focal length of telescope.

Suppose we take a single short-exposure photograph of the object. Then the intensity distribution in the image is given by [13]

$$I(x) = \int_{-\infty}^{\infty} I_o(\alpha)\, S\left(x + \frac{f}{z}\,\alpha\right) d\alpha \qquad (1)$$

where z is the distance of the object from the telescope. Taking the Fourier transform of both sides of (1) yields

$$\tilde{I}(u) = \tilde{I}_o^*\left(\frac{fu}{z}\right)\, \tilde{S}(u) \qquad (2)$$

The quantity $\tilde{S}(u)$ is the instantaneous incoherent optical transfer function of the atmosphere-telescope combination. In the above, as well as succeeding equations, we assume one-dimensional imaging for notational convenience and we restrict ourselves to incoherent imaging.

If we simply average a large number of short exposure images (say of duration 10 ms), we have from (1)

$$\langle I(x)\rangle = \int_{-\infty}^{\infty} I_o(\alpha)\ \langle S\left(x + \frac{f}{z}\alpha\right)\rangle d\alpha \qquad (3)$$

where $\langle S(x)\rangle$ is the average of a large number of instantaneous point spread functions. Similarly, in the frequency domain,

$$\langle \tilde{I}(u)\rangle = \tilde{I}_o^{*}\left(\frac{fu}{z}\right)\langle \tilde{S}(u)\rangle \qquad (4)$$

where $\langle \tilde{S}(u)\rangle$ is the average transfer function. We assume that I_o is constant throughout the averaging process.

We have already pointed out that long-time averaging (or equivalently, time-exposure photography) results in appreciable loss of high-spatial-frequency information. Figure 4 illustrates this phenomenon in terms of the averaged point spread function $\langle S(x)\rangle$. Figure 4a shows the diffraction-limited point spread function corresponding to a 200-inch one-dimensional aperture. Figure 4b shows a typical short-exposure point spread function for the atmosphere-telescope combination. Figure 4c shows a typical long-exposure (or averaged) point spread function. Since

this averaged point spread function is of much greater extent than the diffraction-limited point spread function, the averaged image will be much smoother than the diffraction-limited image and will thus contain appreciably less high-frequency detail.

This phenomenon can also be viewed in the frequency domain as illustrated in Figure 5. Figure 5a shows the optical transfer function of the same aperture as that used for Figure 4. Figure 5b shows a typical short-exposure transfer function for the atmosphere-telescope combination, and Figure 5c shows an averaged transfer function. Comparison of Figure 5a with Figure 5c shows the loss of high spatial frequencies incurred by long-time averaging. For many purposes, this is clearly undesirable.

In speckle interferometry, the squared magnitude of the image Fourier transform is calculated before the averaging is performed. Performing this operation on (2) yields

$$\langle |\tilde{I}(u)|^2 \rangle = \left| \tilde{I}_o^* \left(\frac{fu}{z} \right) \right|^2 \langle |\tilde{S}(u)|^2 \rangle \qquad (5)$$

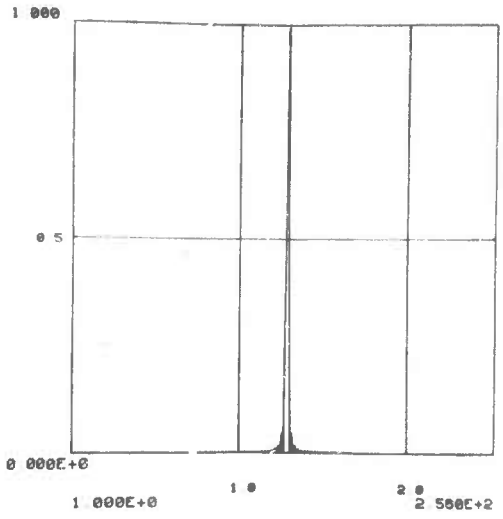

Figure 4a Diffraction-limited point spread function corresponding to a 200-inch one-dimentional aperture.

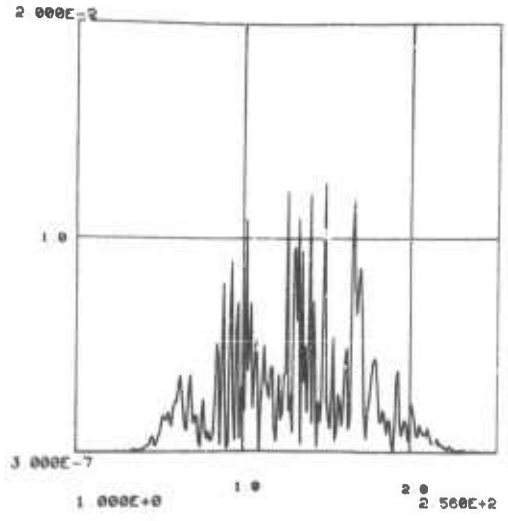

Figure 4b Typical short-exposure point spread function for the atmosphere-telescope combination. Note the speckle-like character.

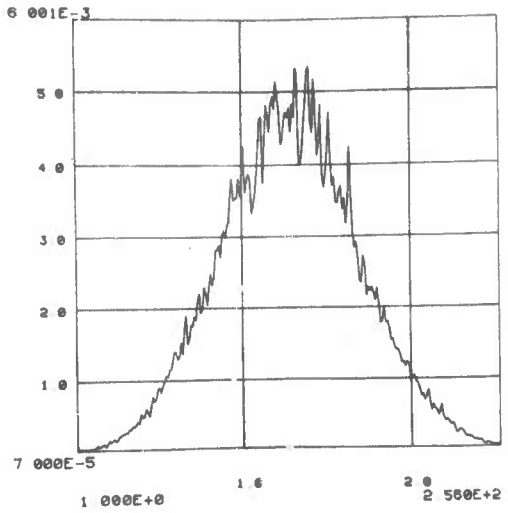

Figure 4c Typical long-exposure (or averaged) point spread function. High spatial frequency detail is attenuated.

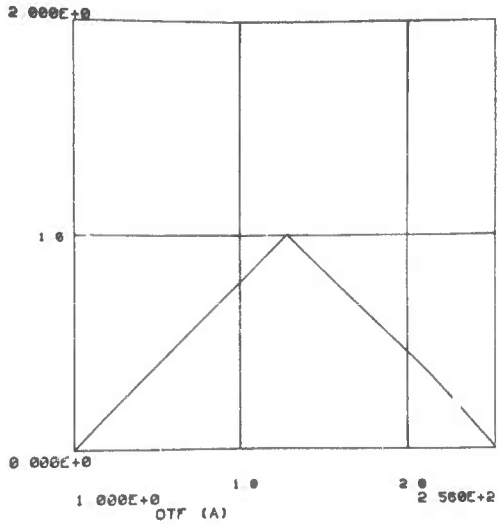

Figure 5a Optical transfer function (OTF) of the same aperture used for Figure 4.

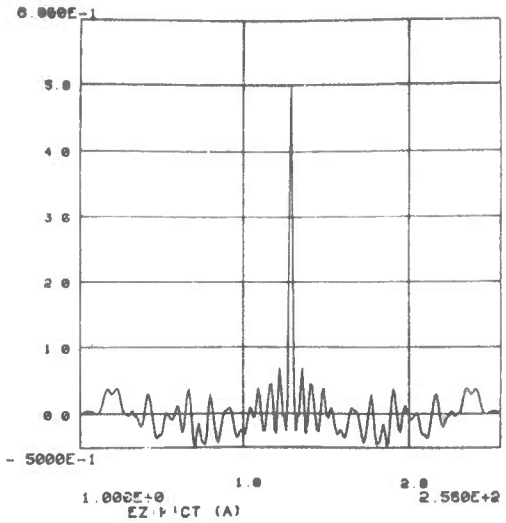

Figure 5b Typical short-exposure transfer function for the atmosphere-telescope combination.

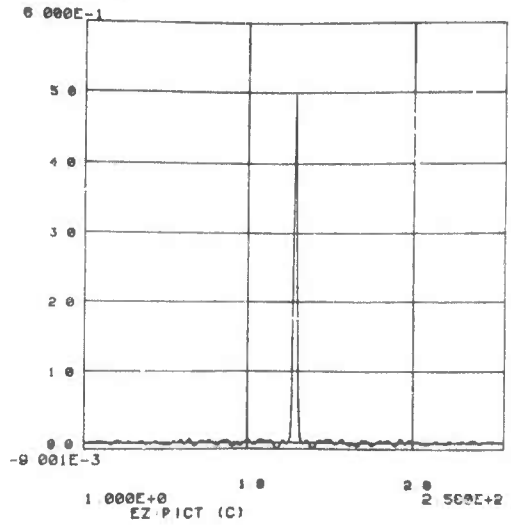

Figure 5c Averaged transfer function. Comparison of Figure 5a with Figure 5c shows the loss of high spatial frequencies incurred by long-time averaging.

Labeyrie has shown experimentally that $\langle|\tilde{S}(u)|^2\rangle$ has useful signal-to-noise ratio out to the diffraction limit of the telescope. Korff [12] and Korff and Miller [11] have provided theoretical explanations and have extended the results to partially coherent illumination. A typical plot of $\langle|\tilde{S}(u)|^2\rangle$ is shown in Figure 13a. Comparison of Figure 13a with Figure 5c illustrates the potential gain in high-frequency information using speckle interferometry.

If we can locate a point source of light not too far removed from the object (but not necessarily in the same isoplanatic patch), we can take a second series of short-exposure photographs of this point source, Fourier transform, square, and average to determine $\langle|\tilde{S}(u)|^2\rangle$. From (5) we will then have

$$\left|\tilde{I}_o^*\left(\frac{fu}{z}\right)\right|^2 = \frac{\langle|\tilde{I}(u)|^2\rangle}{\langle|\tilde{S}(u)|^2\rangle} \qquad (6)$$

In practice, we may wish to modify this procedure to avoid undue amplification of sensor noise, but the basic idea is contained in (6).

Labeyrie and his colleagues [9, 10] have applied speckle interferometry to symmetrical objects with considerable success. Since no phase information is obtained using this technique, however, it will not be effective if applied to arbitrary objects. To overcome this limitation, Knox and Thompson [1] developed a method for calculating the phase of $\tilde{I}_o(u)$ by processing the same set of short-exposure photographs in a somewhat different way. We now turn to a discussion of their procedure.

The effective pupil function of the atmosphere-telescope combination is the product of the telescope pupil function of $A(\hat{x})$ and the instantaneous wavefront $V(\hat{x})$. In order to obtain further results, we must make some assumptions about $V(\hat{x})$. It has been shown that amplitude perturbations caused by atmospheric turbulence are less serious than the effects of phase perturbations. Thus, to a good approximation, we may neglect the amplitude variations and approximate the effects of the atmosphere by a phase perturbation:

$$V(\hat{x}) = e^{i\phi(\hat{x})} \tag{7}$$

We further assume that $\phi(\hat{x})$ is a Gaussian random process with zero mean and autocorrelation function

$$\langle \phi(\hat{x}_1)\ \phi(\hat{x}_2) \rangle = \sigma^2 k^2\ P(\hat{x}_2 - \hat{x}_1) \qquad (8)$$

This assumption implies that we have removed any gross phase shift by centering the photographs before processing.

The instantaneous optical transfer function can now be written as the autocorrelation function of the effective atmosphere-telescope pupil function [13]. Combining this result with the above assumption yields

$$\tilde{S}(u) = \int_{-\infty}^{\infty} e^{i\phi(\hat{x}) - i\ \phi\left(\hat{x}+\frac{f}{k}\ u\right)}\ A(\hat{x})\ A^*\left(\hat{x} + \frac{f}{k}u\right)\ d\hat{x} \qquad (9)$$

Using this equation and the assumed properties of the phase process $\phi(\hat{x})$, we can calculate the properties of $\tilde{S}(u)$.

Although we have already discussed the long-time average transfer function $\langle \tilde{S}(u) \rangle$, it is instructive to calculate this quantity from (9) using the properties of $\phi(\hat{x})$. Averaging both sides of (9) and interchanging orders of expectation and integration on the right side yields

$$\langle \tilde{S}(u) \rangle = \int_{-\infty}^{\infty} \langle e^{i\phi(\hat{x})-i\,\phi\left(\hat{x}+\frac{f}{k}u\right)} \rangle A(\hat{x})\, A^*\left(\hat{x}+\frac{f}{k}u\right) d\hat{x} \qquad (10)$$

The average inside the integral is just the joint characteristic function of the Gaussian random process $\phi(\hat{x})$ and is given by

$$\langle e^{i\phi(\hat{x})-i\,\phi\left(\hat{x}+\frac{f}{k}u\right)} \rangle = e^{-\sigma^2k^2+\sigma^2k^2P\left(\frac{f}{k}u\right)} \qquad (11)$$

Defining

$$K(\hat{x}) = e^{-\sigma^2k^2+\sigma^2k^2P(\hat{x})} \qquad (12)$$

and substituting (11) and (12) into (10) yields

$$\langle \tilde{S}(u) \rangle = K\left(\frac{f}{k}u\right) \int_{-\infty}^{\infty} A(\hat{x})\, A^*\left(\hat{x}+\frac{f}{k}u\right) d\hat{x} \qquad (13)$$

An interesting feature of (13) is the fact that the long-time average transfer function $\langle\tilde{S}(u)\rangle$ appears as the product of two factors, one depending only on the telescope and the other depending only on the atmosphere. The integral appearing in (13) is just the transfer function of the telescope, while $K\left(\frac{f}{k}u\right)$ represents the effect of the turbulent atmosphere. For a high-resolution telescope, K will be many times narrower that the diffraction-limited OTF, and will thus seriously limit the usefulness of $\langle\tilde{S}(u)\rangle$.

The central feature of the Knox-Thompson procedure involves the determination of the phase of the object transform from the autocorrelation function of the image transform. To see how this is done, first recall (2):

$$\tilde{I}(u) = \tilde{I}_0^*\left(\frac{fu}{z}\right)\tilde{S}(u) \tag{2}$$

Calculating the autocorrelation function of both sides of (2) yields

$$\langle\tilde{I}(u_1)\,\tilde{I}^*(u_2)\rangle = \left|\tilde{I}_0^*\left(\frac{fu_1}{z}\right)\right|\left|\tilde{I}_0\left(\frac{fu_2}{z}\right)\right| e^{i\theta\left(\frac{fu_2}{z}\right) - i\theta\left(\frac{fu_1}{z}\right)}$$

$$\cdot\langle\tilde{S}(u_1)\,\tilde{S}^*(u_2)\rangle \tag{14}$$

where

$$\tilde{I}_o(u) = |\tilde{I}_o(u)| \; e^{i\theta(u)} \tag{15}$$

Solving (14) for the phase-difference term yields

$$e^{i\theta\left(\frac{fu_2}{z}\right)-i\theta\left(\frac{fu_1}{z}\right)} = \frac{\tilde{I}(u_1)\,\tilde{I}^*(u_2)>}{|\tilde{I}(u_1)\,\tilde{I}^*(u_2)>|} \cdot \frac{|\tilde{S}(u_1)\,\tilde{S}^*(u_2)>|}{\tilde{S}(u_1)\,\tilde{S}^*(u_2)>} \tag{16}$$

From (16), we see that the difference in the phases of the object transform evaluated at points fu_2/z and fu_1/z can be determined from the autocorrelation functions of the image and the transfer function evaluated at points u_1 and u_2. The phase of $\tilde{I}_o(u)$ can be determined (up to an arbitrary constant) by summing these phase differences.

Knox and Thompson [1] have shown that this procedure is effective provided that the points u_1 and u_2 are sufficiently close together (this statement will be made more precise in the next section). Furthermore, when u_1 and u_2 are sufficiently close together, the phase of $\tilde{S}(u_1)\tilde{S}^*(u_2)>$ is very close to zero and can thus be neglected without serious error. The expression for the phase

difference then reduces to

$$e^{i\theta\left(\frac{fu_2}{z}\right) - i\theta\left(\frac{fu_1}{z}\right)} = \frac{\langle \tilde{I}(u_1)\ \tilde{I}^*(u_2)\rangle}{|\langle \tilde{I}(u_1)\ \tilde{I}^*(u_2)\rangle|} \qquad (17)$$

In summary, the Knox-Thompson procedure involves the following steps:

1. Take a series of short-exposure photographs of the image intensity distribution I(x) through a narrowband optical filter.
2. Take a second series of short-exposure photographs of a point star.
3. Fourier transform the photographs in (1) and (2) to provide sample functions of $\tilde{I}(u)$ and $\tilde{S}(u)$, respectively.
4. Determine the amplitude of the object transform using

$$\left|I_o\left(\frac{fu}{z}\right)\right|^2 = \frac{\langle|\tilde{I}(u)|^2\rangle}{\langle|\tilde{S}(u)|^2\rangle} \qquad (5)$$

5. Determine the phase of the object transform by summing the phase differences obtained from

$$e^{i\theta\left(\frac{fu_2}{z}\right)-i\theta\left(\frac{fu_1}{z}\right)} = \frac{\langle \tilde{I}(u_1)\, \tilde{I}^*(u_2)\rangle}{|\langle \tilde{I}(u_1)\, \tilde{I}^*(u_2)\rangle|} \qquad (17)$$

6. Take the inverse Fourier transform of $\tilde{I}_o(u)$ to find the object distribution $I_o(\alpha)$.

This discussion of the Knox-Thompson technique is by no means complete. The reader interested in more detail should consult [1]. Some additional quantitative detail is provided in Section 3 of the present report as we discuss our experimental procedures and results.

Knox [1] conducted a series of one-dimensional simulations of his procedure in the absence of sensor noise and verified the essential features of its performance. We have conducted a similar series of simulations, and have included the effects of sensor noise. These results are described in Section 3 of this report. Our simulations and

those of others [14] have shown that the Knox-Thompson procedure is very sensitive to sensor noise. These effects can be mitigated somewhat by modifying the procedure to account for noise (e.g., use Wiener filtering to estimate amplitude and phase), but the basic problem remains. A recently-developed alternative to the Knox-Thompson procedure seems to promise greater immunity to sensor noise. This proceduure is briefly described in the next section.

2.5. Sherman's Method

The Knox-Thompson procedure obtains amplitude and phase information separately from two different averages taken over a number of short-time images. The amplitude and phase are then combined and the inverse Fourier transform taken to obtain an estimate of the object intenity distribution. An alternative procedure which treats the object as a whole without explicitly separating amplitude and phase was recently described by Sherman [14].

Sherman's first step is to obtain an integral equation which describes the relationships among the object intensity distribution, the image intensity distribution, the combined effects of the atmospheric turbulence and the telescope, and any imaging noise which might be present. The ultimate objective is to solve this integral equation for the object intensity distribution in terms of measured data and known

statistical parameters.

Sherman proposes the following procedure for obtaining an approximate solution for the object. He first takes the Fourier transform of the images and then calculates the sample covariance matrix of the image transforms. From this quantity he estimates the Cartesian product of the object transform using standard Wiener filtering together with some assumptions about signal and noise statistics. Once an estimate of the Cartesian product is obtained, an estimate of the object itself follows directly from the fact that the only nontrivial eigenvector of the Cartesian product is the object transform itself. If the estimate of the Cartesian product is accurate, the eigenvector associated with the largest eigenvalue should be an accurate estimate of the object transform, although a detailed analysis of the accuracy has not been made.

Sherman has applied this procedure to simulated data with good results. These simulations also indicate that his procedure is less sensitive to noise than is the Knox-Thompson technique.

There are at least three features of the Sherman procedure as described in [14] which should be investigated if that procedure is to become practical and useful. First, and perhaps most important, is the large computational burden associated with the procedure. Second, alternative

techniques for estimating the Cartesian product should be investigated. Third, a more adequate noise analysis using appropriate noise models is needed. With appropriate modifications, Sherman's procedure might well prove quite effective in combating the effects of atmospheric turbulence.

2.6. Real-Time Phase Compensation

An entirely different approach to the reduction of atmospheric turbulence effects is to try to deal with these effects before an image is recorded rather than after. In many circumstances, the effects of atmospheric turbulence can be modeled as a phase perturbation of the received wavefront at the imaging aperture. If this phase perturbation were known, or if an accurate estimate could be obtained, the effects of turbulence could be removed by phase compensation. Several investigators have considered the possibility of implementing real-time phase-compensation systems [15, 16]. In these systems, measurements of optical phase perturbation are made over the objective pupil, either explicitly or implicitly, with the result being used to control a deformable mirror or some other phase-correcting device to compensate for atmospheric turbulence in real time before the image is recorded. Such systems avoid the heavy burden of digital computation which we have described above in connection with the Knox-Thompson and Sherman techniques,

but on the other hand they require more complicated optical equipment and control systems. Furthermore, the accuracy of phase compensation systems is limited by phase-estimation errors due to the finite wavefront-sensor signal-to-noise ratio and by fitting errors due to the finite number of spatial modes restored by the wavefront corrector [17]. Thus, we believe that further extensions and improvements of both phase compensation systems and post-processing techniques of the type we have discussed above should be undertaken.

2.7. The Work of Shapiro

In an interesting series of papers [17-20], Shapiro has established what appears to be a very useful framework for the analysis and design of optical imaging systems which attempt to remove the effects of atmospheric turbulence. In [20], he develops a normal-mode decomposition for the turbulent atmosphere which is similar to that developed previously for free-space imaging by Rushforth and Harris [21] and others. He then applies these results to a study of the ultimate performance limits on imaging through a turbulent atmosphere [17, 19]. In agreement with others, he shows that it is possible in principle to achieve diffraction-limited imaging.

Finally, Shapiro in [18] considers the conditions under

which diffraction-limited imaging may be achieved even when the extent of the object is such that the imaging is no longer isoplanatic; i.e., when the atmospheric point spread function depends upon the position of the point source in the object plane. Since all the work discussed in Sections 2.1 through 2.6 is based on the assumption of isoplanatic imaging, and since in many practical situations the imaging will not be isoplanatic, Shapiro's results are of considerable interest and importance.

In effect, what Shapiro shows is that in many cases of interest, the object can be broken up into distinct isoplanatic elements such that the contribution of each element to the image can be separated from the others. In principle, this enables us to deal with each isplanatic element separately and then put the individual images back together to form the total image. In practice, this procedure may be formidable in its complexity, but it is nevertheless of interest to know that conditions frequently exist under which it is possible in principle.

In the process of obtaining his results on nonisoplanatic imaging, Shapiro [18] develops a model for the atmosphere which may be useful for other purposes. He points out that the problem of optical imaging through a turbulent atmosphere shares many features with the problem of transmitting and receiving information at radio

frequencies through random media such as tropospheric scatter channels. In fact, he shows that the turbulent atmosphere can be modeled as a wide-sense stationary, uncorrelated scatter (WSSUS) channel, which in many respects is the simplest and most useful scatter-channel model [22]. Shapiro modifies and applies the results of [22] to show that when the WSSUS channel is underspread (i.e., the delay-Doppler product is less than one), the effects of the various isoplanatic elements can be separated as described above. This condition frequently occurs in practice.

Thus, we have an example in which a model from random-channel communication theory has proven very useful in the analysis of optical imaging in atmospheric turbulence. In addition, Shapiro has shown that this model can be applied to optical communication systems and to speckle interferometry. We believe that this model, with appropriate modifications, will be very useful in extending the work described in the report.

2.8. Summary

To summarize, we have described in this section several methods which have the potential to achieve high-resolution imaging in the presence of atmospheric turbulence. Considerable work remains before this potential can be realized in a practical system, however. In particular, the

issues of noise sensitivity, computational complexity, and nonisoplanatic imaging must be investigated more carefully. Since sensor noise will be present in any practical system, it is essential that we undertstand how this noise effects various restoration procedures, and that we know how to minimize these effects. The issue of computational complexity may determine whether a technique is hopelessly impractical or can be made operational. Finally, a system which is to be of use in many practical situations must be able to deal with large objects for which the imaging will very likely be nonisoplanatic. We expect the communication theory model described in this section to be useful in these investigations as well as in suggesting alternative or modified procedures with improved performance and efficiency.

3. Simulation of Knox-Thompson Procedure

3.1. Noise Simulation with the Knox-Thompson Technique

A one-dimensional simulation of the Knox-Thompson procedure for restoring atmospheric-turbulence-degraded images was conducted to obtain a better quantitative understanding of its performance in the presence of noise. This work utilized existing models of atmospheric turbulence and of sensor noise. Previous simulations of this technique using data free of sensor noise have led to results in

general agreement with predictions. However, it has been indicated that the technique may be quite sensitive to noise. It was our purpose, therefore, to determine what signal-to-noise ratio was necessary to yield an adequate reconstruction of a degraded image.

3.2. Implementation of the One-Dimensional Knox-Thompson Simulation

This technique uses the assumption that the distortions introduced by the atmospheric turbulence can be described by a Gaussian phase model as discussed in Section 2.4. This model has been used by several other authors to analyze the effects of atmospheric turbulence. It is also assumed that phase distortions are the principal mechanism responsible for the image degradation. The effect of the atmosphere was simulated by randomizing the phase of the pupil function of the imaging system.

An array of Gaussian random noise variables with unit variance was generated. This was filtered by the technique described in Knox [1] to yield correlated noise. The correlation length corresponded to a turbulence cell size of 12 inches across a 200-inch telescope. Figure 6 shows 128 independent arrays of random numbers, each array being shown left to right and successive arrays top to bottom. The arrays were obtained from a random process which was zero

mean, unit variance, and uncorrelated. The impulse response of the digital filter is shown in Figure 7. The correlated noise is shown in Figure 8. The correlated noise arrays were used as the phase $\phi_k(x)$ and the aperture function $A(x)$ shown in Figure 9 as the magnitude to form the complex pupil functions. The pupil functions were inverse Fourier transformed and the squared magnitude formed yielding a set of 128 different point spread functions $S_k(x)$, $k = 1, 2, \ldots, 128$.

$$S_k(x) = \left| F^{-1}\left(A(x)\, e^{i\phi_k(x)} \right) \right|^2$$

The array of 128 point spread functions is shown in Figure 10. Each function is shown left to right and successive functions are shown top to bottom.

A second set of 128 point spread functions was generated in a similar manner. These were convolved with the double star image shown in Figure 11 to yield a set of 128 blurred images $I_k(x)$. One typical blurred image is shown in Figure 12. These blurred double star images were used as the input for the Knox-Thompson restoration.

Let a tilde represent the Fourier transform of a function. Thus $\tilde{S}$ and $\tilde{I}$ represent the Fourier transform of S

and I, respectively.

Both Labeyrie's technique and the Knox-Thompson technique estimate the magnitude of the Fourier transform of the image $|\tilde{I}_o|$ in the same manner. The first step is to average the magnitude squared of the Fourier transform of the point spread functions; i.e., to form the quantity $\langle|\tilde{S}_k(u)|^2\rangle$, shown in Figure 13a. The brackets represent an average over the 128 spread functions. In a similar manner, the average of the magnitude squared of the Fourier transform of blurred images is formed, i.e., $\langle|\tilde{I}_k(u)|^2\rangle$. This quantity is shown in Figure 13b. Then the estimate of the magnitude of the image transform $|\tilde{I}_o|$ is given by

$$|\tilde{I}_o| = \left(\frac{\langle|\tilde{I}_k(u)|^2\rangle}{\langle|\tilde{S}_k(u)|^2\rangle}\right)^{1/2}$$

If the inverse Fourier transform of $|\tilde{I}_o|$ is taken using zero phase, the autocorrelation of the object is obtained. This is the Labeyrie result, and an example is shown in Figure 13c.

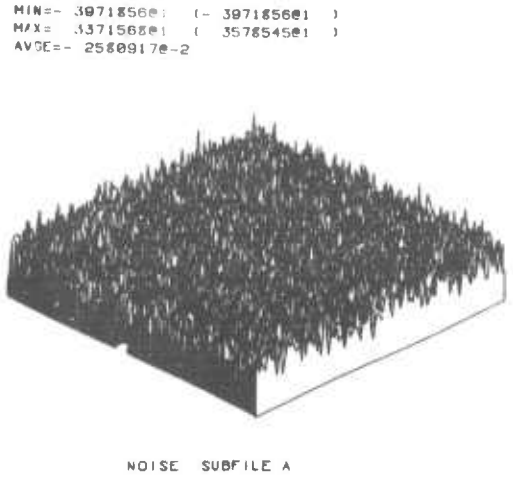

Figure 6 Isometric plot of 128 independent arrays of random numbers. Each array is shown going upwards to the right and successive arrays upward to the left. The arrays were obtained from a zero mena, unit variance and uncorrelated random process.

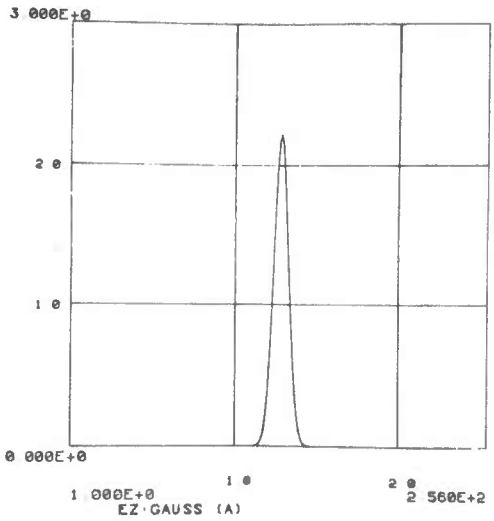

Figure 7 Impluse response of the gaussian filter used to filter the uncorrelated noise of Figure 6 to produce the correlated noise in Figure 8.

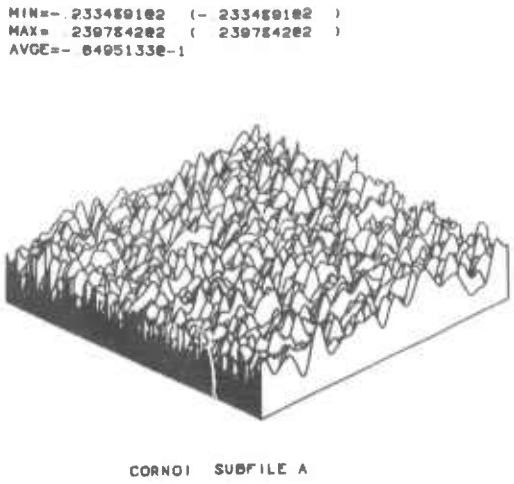

Figure 8 Correlation noise used as the phase $\phi(x)$ of the complete pupil function.

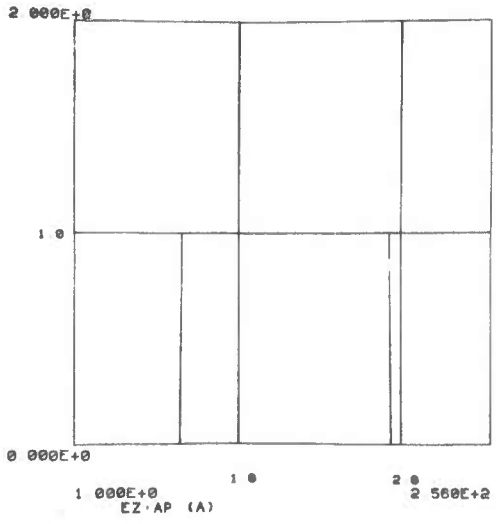

Figure 9 Telescope aperture function A(x), which was used as the magnitude of the complex pupil function.

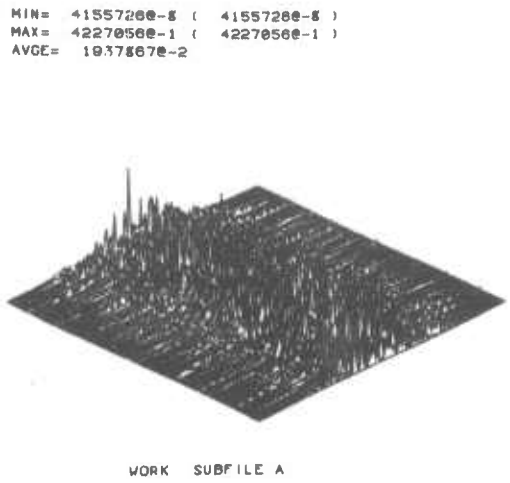

Figure 10 Isometric plot of 128 point spread functions. These were formed by inverse Fourier transforming the complex pupil functions and then obtaining the squared magnitude.

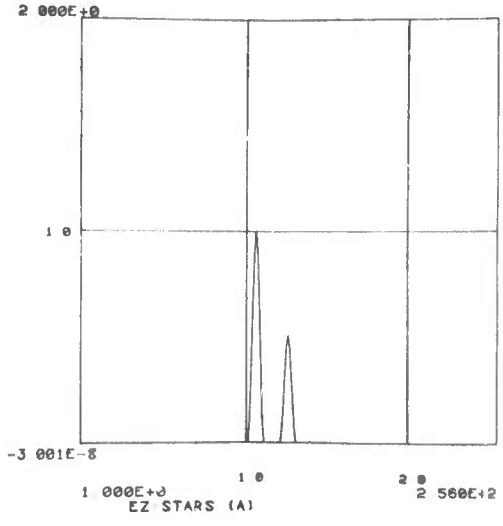

Figure 11 One-dimensional intensity pattern of a double star.

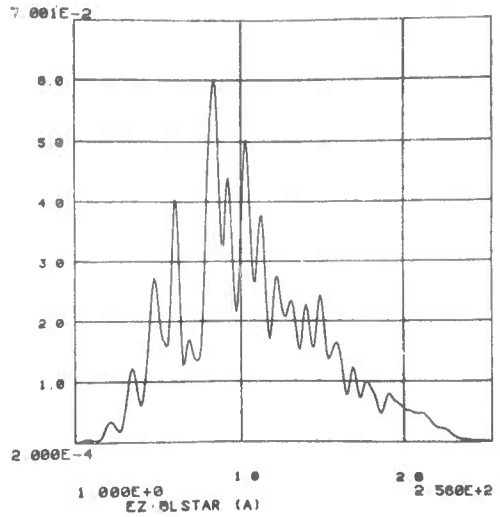

Figure 12 Example of a blurred double star obtained by convolving the image in Figure 11 with one of the point spread functions of Figure 10.

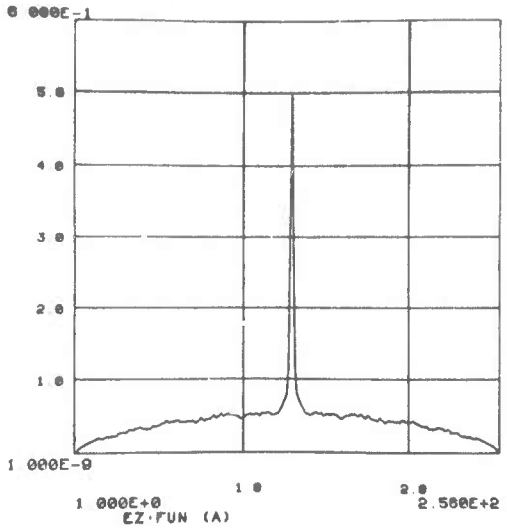

Figure 13a The average of the magnitude squared of the Fourier transform of the point spread functions shown in Figure 10.

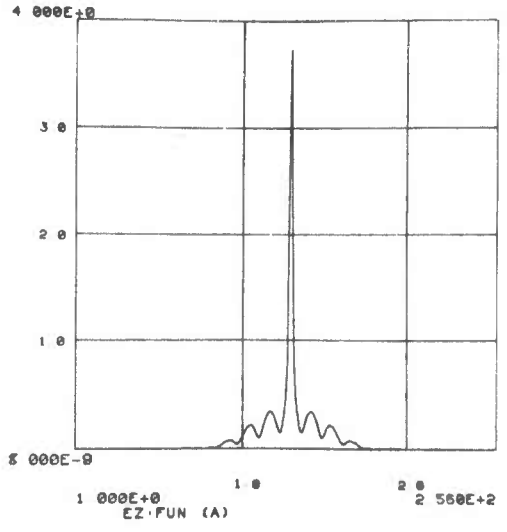

Figure 13b Average of the magnitude squared of the Fourier transform of the blurred double stars, one of which is shown in Figure 12.

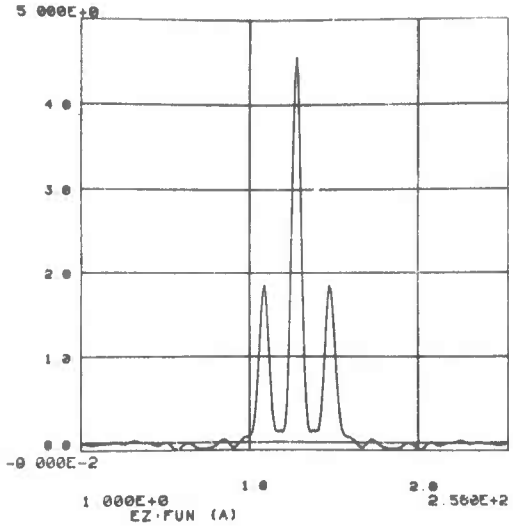

Figure 13c Labeyrie reconstruction of a double star obtained by dividing Figure 13b by Figure 13a, taking the square root, and then inverse transforming using zero phase.

Knox and Thompson have gone one step further and provide a method to obtain phase information. They have shown that differential phase may be obtained from the quantity $\langle \tilde{I}(u_1)\, \tilde{I}^*(u_2) \rangle$, where u_1 and u_2 are spatial frequencies differing from each other by a small amount. In our case, they differed by one sample in the frequency domain.

The technique is to average the differential phase to produce an object phase distribution. This is combined with the amplitude of the object transform previously obtained and inverse transformed. The result is the object intensity distribution with spatial frequencies present out to the diffraction limit of the telescope. This technique thus offers the potential to improve the resolution of arbitrary intensity distributions.

A detailed description of the calculation of the phase will now be given. The first step is to shift each point spread function and each blurred image such that they have no movement about their center. This step is to remove large linear phase shifts and has been shown [1] to greatly improve the contrast of the resultant reconstruction.

The centered blurred images are then Fourier transformed. The transforms are shifted one step and complex conjugated. The product of the original and shifted transforms is formed and the average over the 128 images is

taken. In mathematical terms, the quantity $\langle \tilde{I}(u_1)\ \tilde{I}^*(u_2) \rangle$ is obtained. In a similar manner, the quantity $\langle \tilde{S}(u_1)\ \tilde{S}^*(u_2) \rangle$ is obtained. The phase of the quantity

$$\frac{\langle I(u_1)\ I^*(u_2) \rangle}{\langle S(u_1)\ S^*(u_2) \rangle}$$

is just the phase difference between adjacent points u_1 and u_2 in the object transform array.

The phase of the object transform, ϕ_o, is obtained by summing the phase differences from the origin outward (see Figure 14a). This phase is combined with the quantity $|\tilde{I}_o|$ and the inverse Fourier transform taken:

$$I_o = F^{-1} \left\{ |\tilde{I}_o|\ e^{i\phi_o} \right\}$$

The resultant deblurred image is shown in Figure 14b.

This result, with no sensor noise, shows a good representation of the original double star. The orientation of the large star relative to the small star was reversed as predicted.

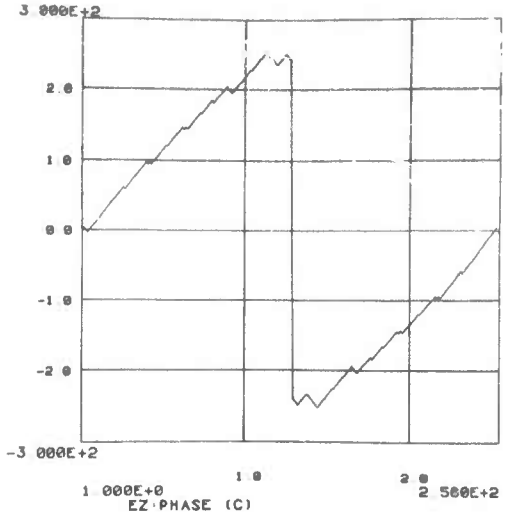

Figure 14a Phase obtained by the Knox-Thompson technique. DC phase is at left.

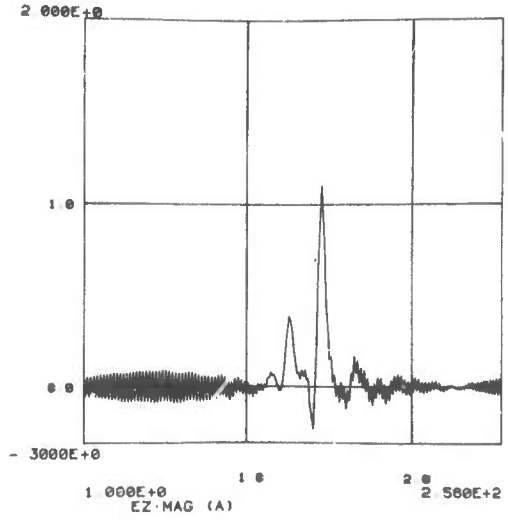

Figure 14b Double star image obtained with the Knox-Thompson technique.

3.3. Effects of Sensor Noise

In order to obtain a more quantitative understanding of the effects of sensor noise on the Knox-Thompson procedure, we performed a series of simulations in which controlled amounts of sensor noise were present. The Knox-Thompson restoration procedure was then applied to the noisy data, and the resulting restored images were compared with those obtained in the absence of sensor noise.

The sensor-noise model we used to generate our noisy images is based on a semiclassical approach to photon detection [23]. In this model, the number of photoelectrons released from a small region of area ΔA centered at a point (x, y) in the image plane is taken to be a Poisson random variable with parameter

$$\lambda = \frac{\eta\, I(x, y)}{h\, \bar{\nu}}\, \tau\, \Delta A \tag{18}$$

In this expression, η is the quantum efficiency of the photodetector, h is Planck's constant, $\bar{\nu}$ is the mean optical frequency, τ is the integration (exposure) time, and $I(x, y)$ is the image intensity at point (x, y).

It is well known that both the mean and the variance of the above Poisson random variable are equal to λ. Thus, the mean photoelectron current produced by a detector at point (x, y) is proportional to the image intensity at that point, but so are the fluctuations in that current. If we define the signal-to-noise ratio of the image at point (x, y) as the ratio of the square of the mean current to the variance of the current, we find that this signal-to-noise ratio is just λ. We can vary λ, and therefore the signal-to-noise ratio, by varying τ (the integration time) or ΔA (the area over which the image is averaged).

An additional simplification results if we assume that the number of photoelectrons is large. In this case, the photoelectric current will be approximately Gaussian with a mean equal to its variance. We made this assumption in the simulations which we performed. The signal-to-noise ratio which we ascribe to a given simulated image is simply the maximum of the point-by-point ratios of squared mean to variance as described above. This definition is somewhat arbitrary, of course, but this is not a serious problem since we are interested primarily in relative performance as we vary the noise.

Results of the Knox-Thompson restoration procedure for various signal-to-noise ratios as defined above are shown in Figures 15a through 15j. Examples of Labeyrie

speckle-interferometry reconstructions are shown in Figures 16a through 16d.

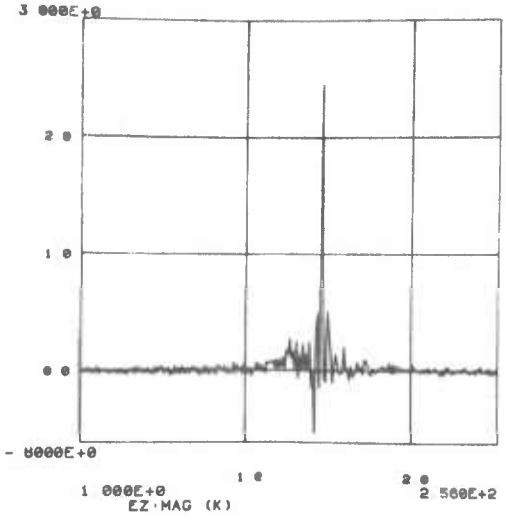

SNR = 5db

(a)

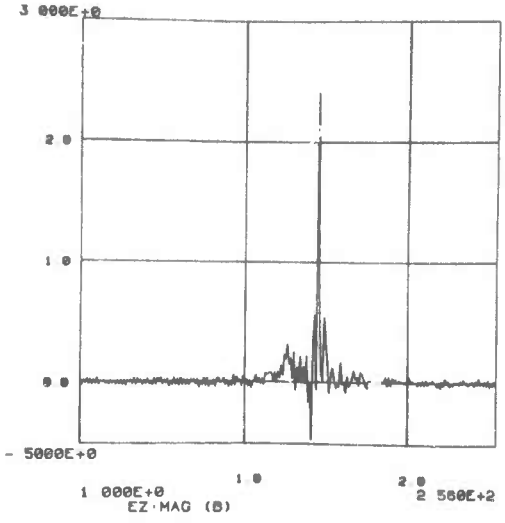

SNR = 10db

(b)

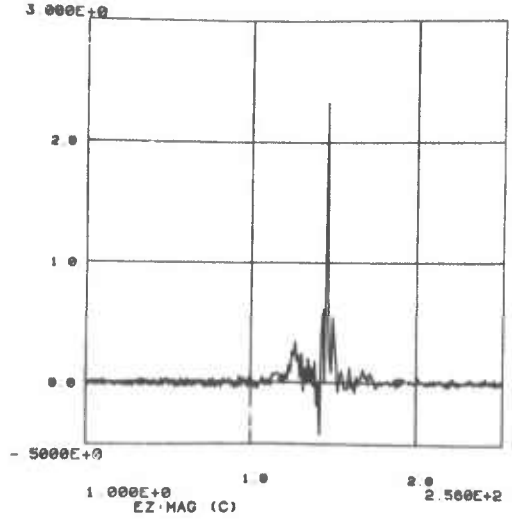

SNR = 13db

(c)

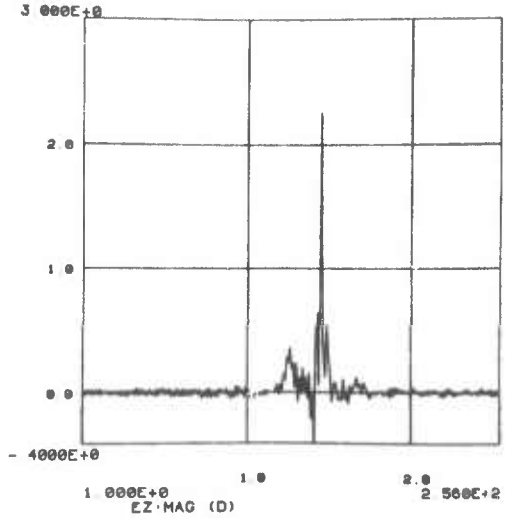

SNR = 15db

(d)

Figure 15 Knox-Thompson restorations of a double star signal to noise ratio (SNR) as indicated under each plot.

Figure 15 Continued

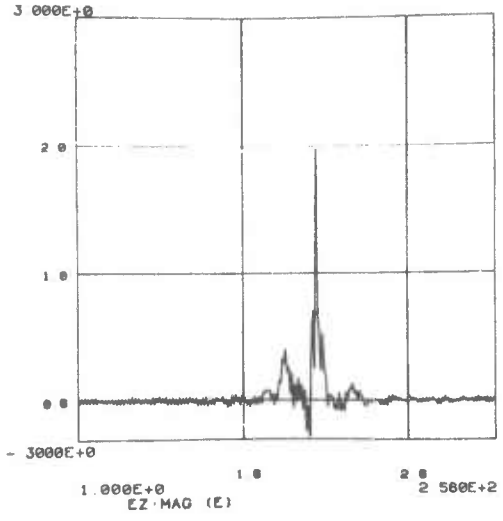

SNR = 20db

(e)

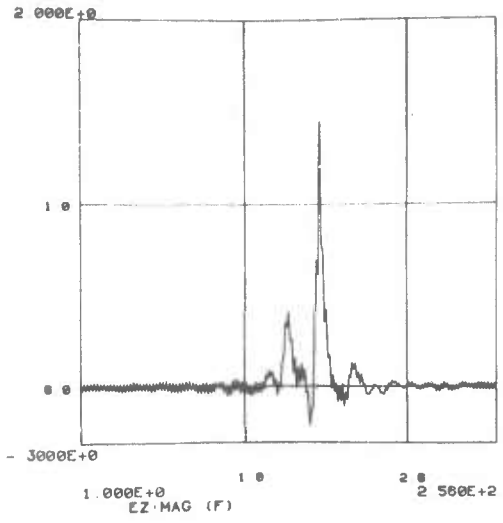

SNR = 30db

(f)

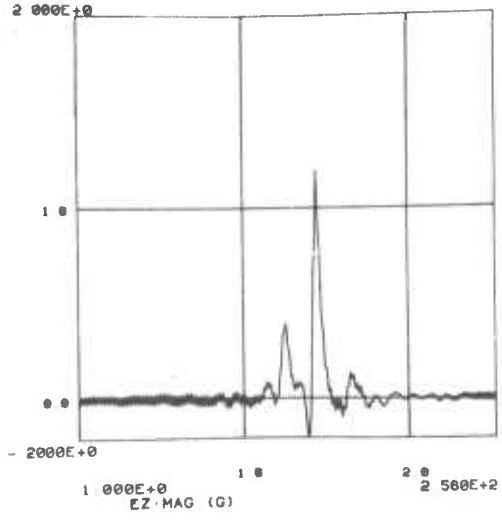

SNR = 40db

(g)

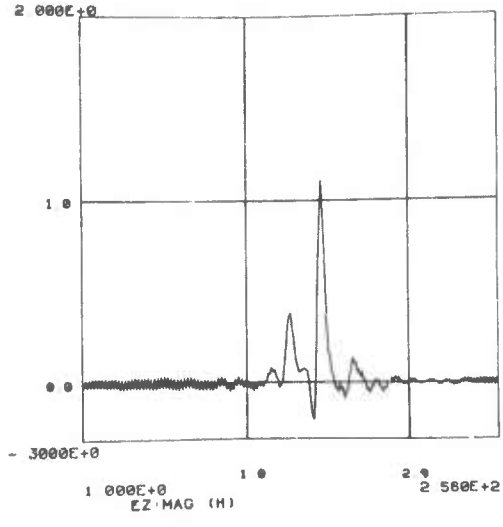

SNR = 50db

(h)

Figure 15 Continued

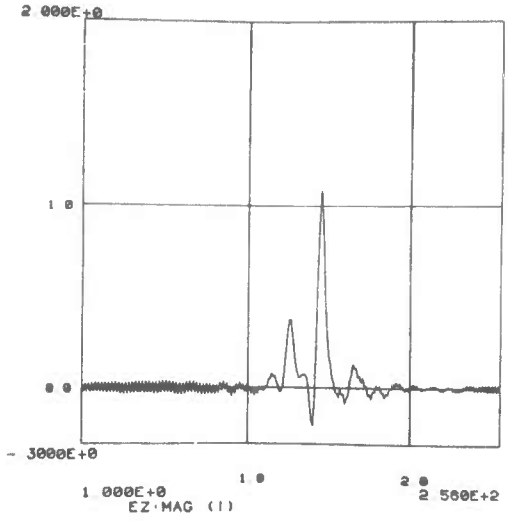

SNR = 60db

(i)

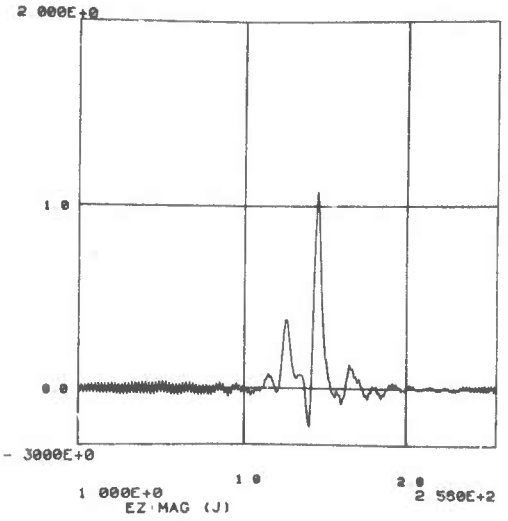

SNR = 70db

(j)

Figure 15

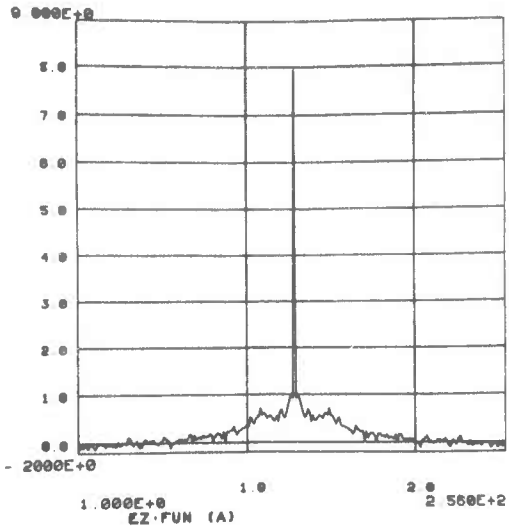

SNR - 5db

(a)

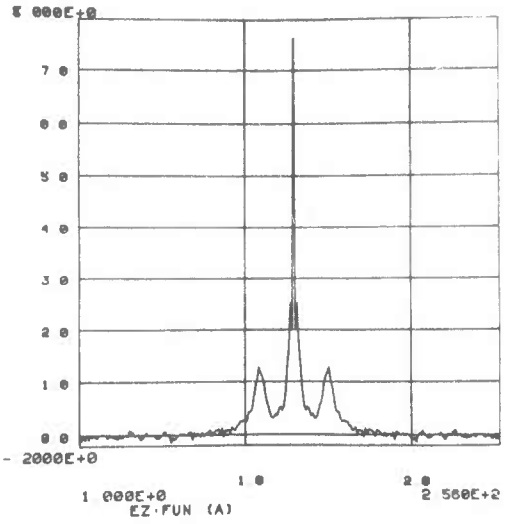

SNR = 13db

(b)

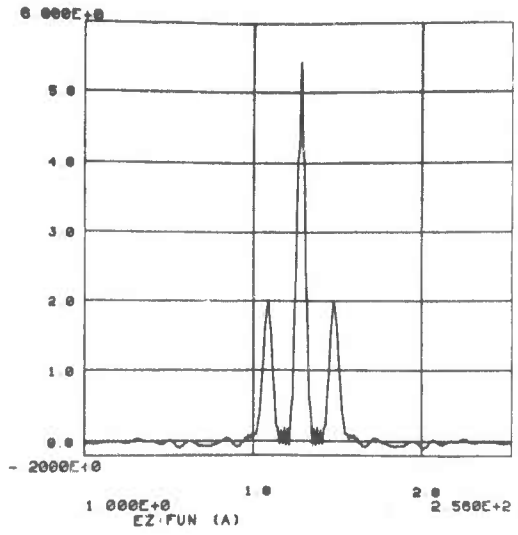

SNR = 40db

(c)

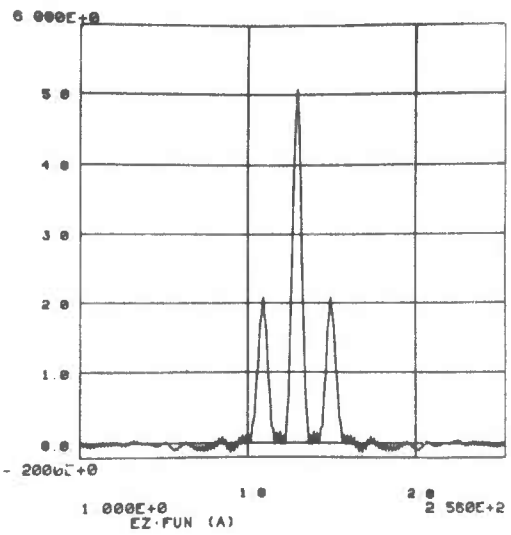

SNR = 70db

(d)

Figure 16 Labeyrie reconstruction of blurred double star with specified signal to noise ratios.

3.4. Conclusions on the Effects of Noise

Figures 15a through 15j show the results of the Knox-Thompson reconstruction of the blurred double star image with sensor noise. The signal-to-noise ratio (SNR) varied from 5dB up to 70dB. Figure 15a has a SNR of 5dB. There is no indication of a second star. Figures 15b with a 10dB SNR could possibly indicate a second star. Figure 15c through 15f with SNRs varying from 13dB to 30dB definitely show a second star, but the magnitudes of both the larger and smaller stars are distorted. It is not until a 40dB SNR shown in Figure 15g that the magnitudes approach their true value. For higher SNRs (50dB to 70dB), there appears to be little improvement in the reconstructed images.

For SNRs of 30dB and greater, a third "star" appears to the right of the main star. Its magnitude is about one-tenth the magnitude of the main star. Increasing the SNR did not reduce its presence, but rather made it more distinct.

Figures 16a through 16d show Labeyrie reconstructions with SNRs varying from 5dB to 70dB. Figure 16a with a 5dB SNR shows little or no presence of a second star. Figure 16b with a SNR of 13dB definitely shows a second star, but the magnitudes are incorrect. For a SNR of 40dB, the magnitudes are correct. Increasing the SNR to 70dB shows

little improvement.

REFERENCES

[1] Keith T. Knox, "Diffraction-Limited Imaging with Astronomical Telescopes", Ph.D. dissertation, Institute of Optics, University of Rochester, New York, 1975.

[2] R.E. Hufnagel and N.R. Stanley, "Modulation Transfer Function Associated with Image Transmission through Turbulent Media", Journal of the Optical Society of America, January 1964, pp. 52-61.

[3] D.L. Fried, "Optical Resolution Through a Randomly Inhomogeneous Medium for Very Long and Very Short Exposures", Journal of the Optical Society of America, Vol. 56, 1966, p. 1372.

[4] A.A. Michelson, "On the Application of Interference Methods to Astronomical Measurements", Astrophysical Journal, Vol. 51, 1920, p. 257.

[5] J.A. Anderson, "Application of Michelson's Interferometer Method to the Measurement of Close Double Stars", Astrophysical Journal, Vol. 51, 1920, p. 263.

[6] R. Hanbury-Brown and R.Q. Twiss, "Interferometry of the Intensity Fluctuations in Light", Proceedings of the Royal Society A, Vol. 242, 1957, p. 300.

[7] R.H. MacPhie, "The Compound Intensity Interferometer", IEEE Transactions, Vol. AP-14, 1966, p. 369.

[8] R.H. MacPhie, "The Quasi-Linear Intensity Interferometer", IEEE Transactions, Vol. AP-20, 1972, p. 755.

[9] A. Labeyrie, "Attainment of Diffraction Limited Resolution in Large Telescopes by Fourier Analyzing Speckle Patterns in Star Images", Astronomy and Astrophysics, Vol. 6, 1970, pp. 85-87.

[10] D.Y. Gezari, A. Labeyrie, and R.V. Stachnik, "Speckle Interferometry: Diffraction-Limited Measurements of Nine Stars with the 200-Inch Telescope", Astrophysical Journal, April 1, 1972, pp. L1-L5.

[11] M.G. Miller and D. Korff, "Resolution of Partially Coherent Objects by Use of Speckle Interferometry", Journal of the Optical Society of America, February 1974, pp. 155-161.

[12] D. Korff, "Analysis of a Method for Obtaining Near-Diffraction-Limited Information in the Presence of Atmospheric Turbulence", Journal of the Optical Society of America, August 1973, pp. 971-980.

[13] J.W. Goodman, "Introduction to Fourier Optics", McGraw-Hill Book Company, Inc., New York, 1968.

[14] J.W. Sherman, "A Posteriori Restoration of Atmospherically Degraded Images Using Multiframe Imagery", Topical Meeting on Image Processing, Optical Society of America, Pacific Grove, California, February 24-26, 1976.

[15] R.A. Muller and A. Buffington, "Real-Time Correction of Atmospherically Degraded Telescope Images through Image Sharpening", Journal of the Optical Society of America, September 1974, pp. 1200-1210.

[16] J.W. Hardy, J. Feinleib, and J.C. Wyant, "Real-Time Phase Correction of Optical Imaging Systems", OSA Topical Meeting on Optical Propagation Through Turbulence, Boulder, Colorado, July 1974.

[17] J.H. Shapiro, "Propagation-Medium Limitations on Phase-Compensated Atmospheric Imaging", Journal of the Optical Society of America, May 1976.

[18] J.H. Shapiro, "Diffraction-Limited Atmospheric Imaging of Extended Objects", Journal of the Optical Society of America, May 1976, pp. 469-477.

[19] J.H. Shapiro, "Optimum Adaptive Imaging through Atmospheric Turbulence", Applied Optics, November 1974, pp. 2609-2613.

[20] J.H. Shapiro, "Normal-Mode Approach to Wave Propagation in the Turbulent Atmosphere", Applied Optics, November 1974, pp. 2614-2619.

[21] C.K. Rushforth and R.W. Harris, "Restoration, Resolution, and Noise", Journal of the Optical Society of America, April 1968, pp. 539-545.

[22] R.S. Kennedy, "On the Atmosphere as an Optical Communication Channel", IEEE Transactions on Information Theory, September 1968, pp. 716-725.

[23] J.W. Goodman and J.F. Belsher, "Fundamental Limitations in Linear Invariant Restoration of Atmospherically Degraded Images", Technical Note, Department of Electrical Engineering, Stanford University, Stanford, California, 1975.

SECTION III

IMAGE UNDERSTANDING

Martin E. Newell

1. Introduction

The work done from mid 1976 to March 1977 in the area of Image Understanding continued to produce the tools necessary for implementing an analysis by synthesis image processing facility. The analysis of images is carried out for the purposes of automatic recognition of previously known objects, or for synthesizing models of previously unknown objects. The basic notion is that such analyses can benefit greatly if carried out in conjunction with three-dimensional models of the objects in the scene.

Given modelling and image synthesis facilities of sufficiently advanced capability, analysis of such scenes can be carried out using an analysis by synthesis approach. Based on some hypothesis about the objects in the scene, and their orientations with respect to the camera, a synthetic image can be created and compared with the actual image. The model of the objects in the scene is then modified based on differences between these two images, and the cycle repeated.

2. Main Subproblems

Four main problem areas can be identified in attempting to develop such a system:

1. Abstraction of perceptually relevant information from images, for use in making comparisons between real and synthesized images.
2. Generalized correlation in both 2-D and 3-D for the purposes of finding the best fit between the real image and a synthetic one.
3. Synthesis of high fidelity images, capable of reproducing the perceptually important characteristics of real images.
4. Advanced modelling systems, for storing and manipulating a wide variety of object representations.

The progress that has been made in these four areas is described below.

2.1. Abstraction of Perceptually Important Information

In any Image Understanding task, it is necessary that the features of an image relevant to the perception of the subject of the image be readily identified. The most obvious and commonly used features of images are high contrast boundaries, or edges. Several heuristic techniques for detecting edges in images have been developed, some more successful than others (Hueckel (1971)[1], Griffith

(1973)[2]).

All of these techniques suffer to a greater or lesser degree from the effects of noise, or of widely varying local contrast in the image. In an attempt to provide better techniques for abstracting perceptually important information from images, a new analysis method has been developed. This new analysis involves the creation of a transform of the image, called a Mandalagram.

A Mandalagram is the two-dimensional analog of a short time spectogram, used in the analysis of acoustic waveforms. A Mandalagram is constructed by fragmenting the given image into an array of small overlapping areas, typically 8 x 8 picture elements. The two-dimensional Fourier transform of each of these areas is calculated. The Mandalagram is then constructed by taking corresponding components from all of the small area Fourier transforms, and assembling them into a mosaic of images, one image for each component of the transforms, each component of the transform being located in its image in a position corresponding to the position of its small area in the original image.

The mathematics underlying the Mandala transform is described in Kajiya (1976)[4].

2.2. Generalized Correlation

The ultimate goal of this part of the work is to provide correlation techniques to enable the subject of a given image to be matched with one of a set of given prototype objects. The prototype objects may be parameterized with continuously variable parameters, in which case the values of the parameters to give a best fit to the imaged object are sought.

The problem has been approached in several steps of increasing difficulty. In the first step the given image and a prototype image are assumed to differ with four degrees of freedom, namely x position, y position, rotation in the plane of the image, and scale. Furthermore, the given image is assumed to show a single object against a black background.

This problem can be approached in two ways - as a linear programming problem in four variables, or by using transform techniques to compute the correlation function. Two versions of the latter are used here.

By way of review, two images which differ in only x position and y position can be cross correlated using the two-dimensional Fourier transform. Let the given and prototype images be $f(x,y)$ and $g(x,y)$ respectively. Then the correlation of the two can be computed as the two-dimensional convolution:

$$c(x,y) = f(x,y) * g(-x,-y) = \langle f,g \rangle$$

In order to find the best fit between the two images, correlation coefficients, $\bar{c}(x,y)$, are needed:

$$\bar{c}(x,y) = \frac{\langle f,g \rangle}{\sqrt{\langle f,f \rangle \langle g,g \rangle}}$$

The juxtaposition of the two images giving a best fit is given by the coordinates, (x,y) for which $\bar{c}(x,y)$ has a maximum. Morever, the magnitude of the maximum gives an indication of the closeness of the fit, unity indicating a perfect fit.

Whereas higher dimensional versions of the above technique can, in principle, be applied in the four degrees of freedom case, the required computational time and space become wholly excessive. Consequently, two other approaches which avoid such problems have been developed. The first of these, called the Transform Method, transforms the images into a space where the effects of translation can be separated from the effects of scale and rotation. This allows the problem to be considered as two two-dimensional correlations in series. The second alternative approach called the Centroid Method, is similar in intent, is computationally more efficient, but more prone to inaccuracies in the presence of noise.

The Transform Method

Consider the Fourier transforms of two images $f(x,y)$ and $g(x,y)$:

$$F(w_1,w_2) = F(f(x,y))$$

$$G(w_1,w_2) = F(g(x,y))$$

Assume that f and g differ by a translation of $(\Delta x,\Delta y)$, a rotation in the x,y plane by α, and a scale s, i.e.

$$g(x,y) = f(\text{trans}(\Delta x,\Delta y)\text{rotate}(\alpha)\text{scale}(s)(x,y))$$

then it may be shown that:

$$G(w_1,w_2) = \frac{1}{|s|} F(\text{rotate})(\alpha)\text{scale}(1/s)(x,y))e^{j(w_1\Delta x+w_2\Delta y)}$$

From this expression we see that the transform is rotated exactly the same as the image, it is scaled by the inverse of the image scale, and its phase is altered by the addition

of planar phase. That is, all three geometric transformations affect phase, but only rotation and scale affect magnitude, in the manner shown. Consequently, the differences in the magnitudes of the two Fourier transforms are entirely due to rotation and scale. Moreover, the magnitudes will themselves differ by the same rotation and inverse scale as do the images.

By working with the magnitudes of the Fourier transforms, the problem is reduced to one having only two degrees of freedom, rotation and scale. This can be solved as a two-dimensional correlation. However, it is first necessary to change the coordinate system to radial coordinates, (r,θ) with exponential radial distances, i.e.

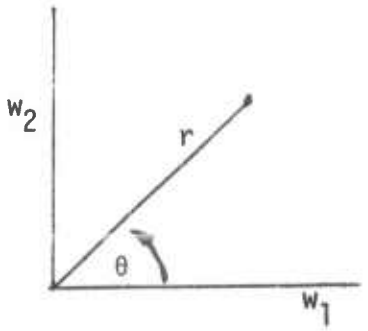

$$r = \ln \sqrt{w_1^2 + w_2^2}$$

$$\theta = \tan^{-1}\left(\frac{w_2}{w_1}\right)$$

$$|F(w_1,w_2)| \rightarrow \phi(r,\theta)$$

$$|G(w_1,w_2)| \rightarrow \gamma(r,\theta)$$

Consider the effect of scaling by 1/s in w_1 and w_2

$$|F(w_1/s, w_2/s)| \rightarrow \phi(r_s, \phi_s)$$

where

$$r_s = \ln \sqrt{(w_1/s)^2 + (w_2/s)^2}$$

$$= \ln \sqrt{(w_1^2 + w_2^2)\,(1/s)^2}$$

$$r_s = \ln \sqrt{w_1^2 + w_2^2} - \ln s$$

$$= r - \ln s$$

and

$$\theta_s = \tan^{-1} \frac{w_1/s}{w_2/s}$$

$$= \tan^{-1} \frac{w_1}{w_2}$$

$$= \theta$$

i.e. A scale by (1/s) in (w_1,w_2) space is mapped into a translation by -ln s in the r coordinate of (r,θ) space. It can be shown that a rotation by α in (w_1,w_2) space maps into a translation by α in the θ coordinate of (r,θ) space. This enables us to again use two-dimensional Fourier transform techniques to cross correlate the two functions. However, the functions $\phi(r,\theta)$ and $\gamma(r,\theta)$ must first be compensated for the distortion of area that occurs in the coordinate transformation. This step is necessary because the correlation is essentially an integral with respect to r and θ and is therefore affected by changes of metric. This compensation can be achieved by multiplying both $\phi(r,\theta)$ and $\gamma(r,\theta)$ by e^{r^2} which is the determinant of the Jacobian of the transformation, i.e.

$$\phi'(r,\theta) = e^{r^2}\phi(r,\theta)$$

$$\gamma'(r,\theta) = e^{r^2}\gamma(r,\theta)$$

As a result of this correlation, we find values of r and θ which give a maximum in the function:

$$\frac{\langle d',j'\rangle}{\sqrt{\langle d',d'\rangle\langle j',j'\rangle}}$$

Let these values be R and θ. Then the scale factor giving a best fit in the original image is

$$\text{Scale} = e^{R}$$

and the rotation is θ.

Having determined the rotation and scale, it remains to find the displacements in x and y. This can be done by applying the scale and rotation found above to the prototype image, then carrying out a two-dimensional cross correlation with the given image to find the displacements. A one-dimensional version of this method is shown in Figure 1.

The transform method for determining the four degrees of freedom may be used in the presence of limited amplitude high frequency noise. Using this method the goal of correlating images with templates in 2-D can be achieved.

The Centroid Method

For the given subproblem of determining the position, size, and orientation of a given image against a black background, another transform technique exists. This involves finding the centroid, C_I, of the given image and C_T of the template. The position of the centroid is

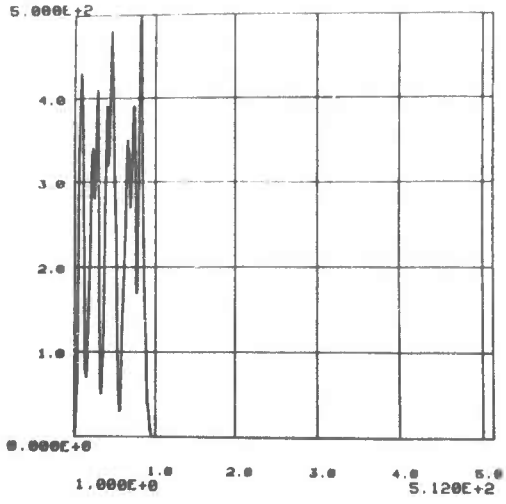

(a)

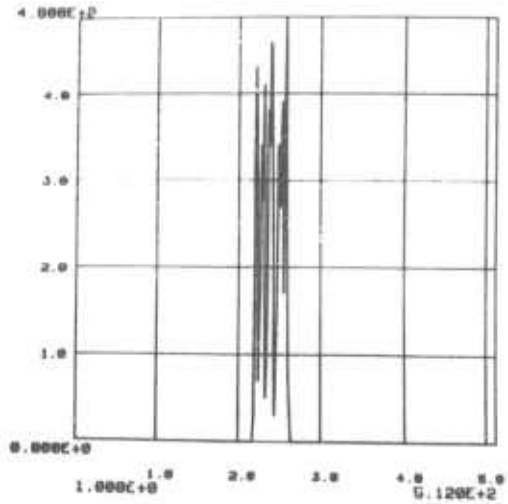

(b)

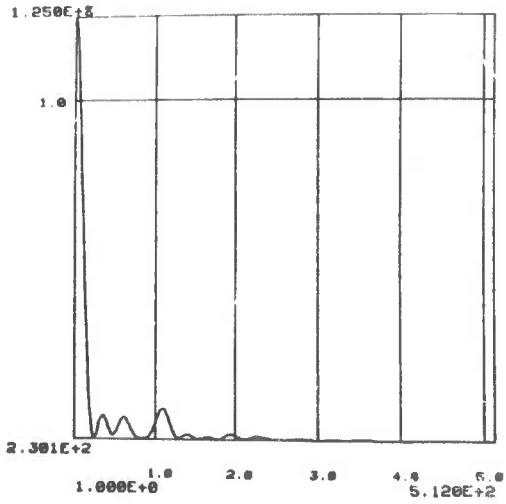

(c)

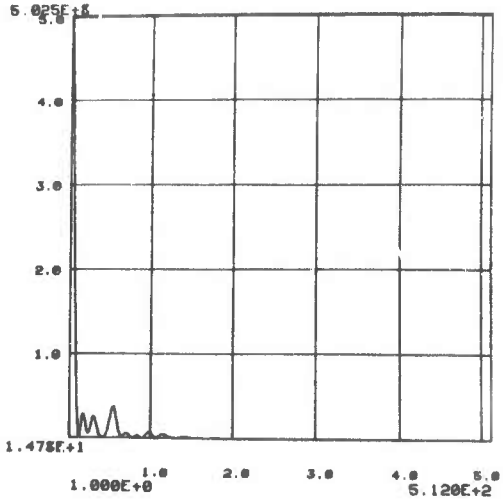

(d)

Figure 1. Generalized correlation in 1-D.

(continued)

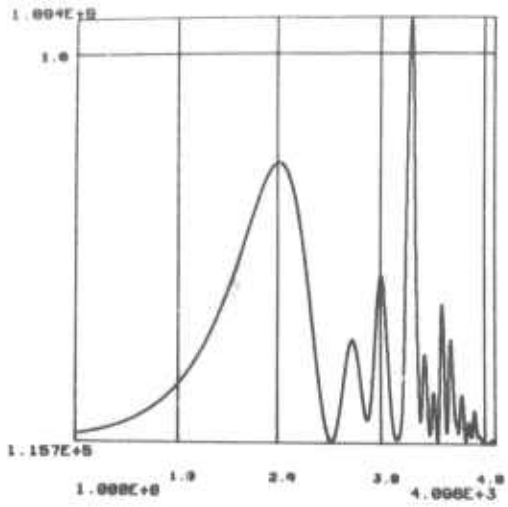

(e)

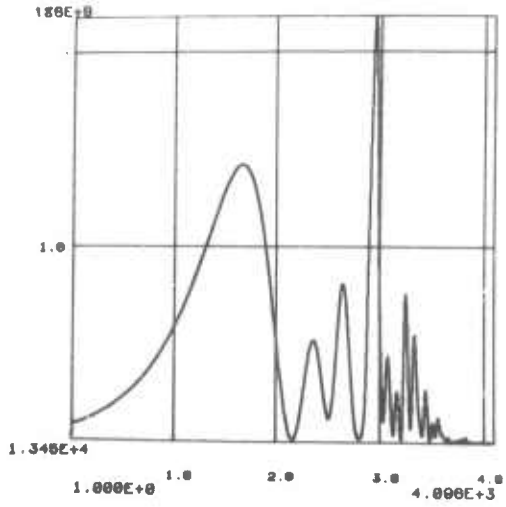

(f)

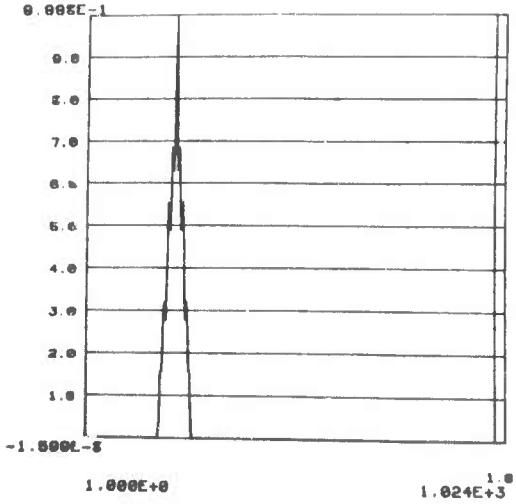

(g)

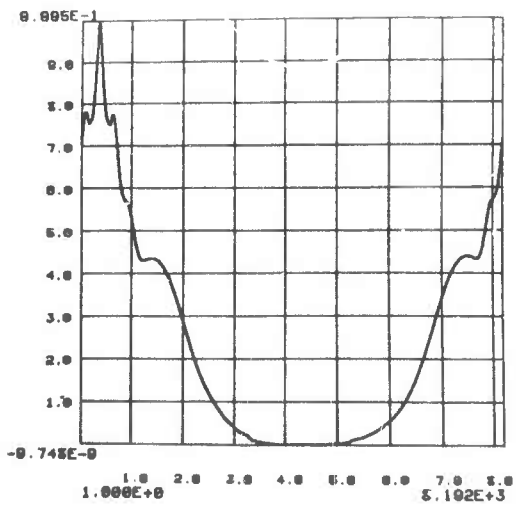

(h)

Figure 1. Generalized correlation in 1-D.

independent of rotation and scale about the centroid. Therefore the difference in the positions of the centroids C_I and C_T will be precisely the translation (Δx Δy) between the given image and the template.

To find the rotation and scale, we translate both the image and template so as to bring their centroids to the origin of coordinates. From here the problem is the same as in the Transform Method, except that now we are in image space instead of in magnitude Fourier transform space.

Since the computation of centroid is extremely simple, the Centroid Method is faster than the Transform Method. However, it is probably more prone to error in the presence of noise in the image (Figure 2).

2.3. High Fidelity Synthetic Images

The basic notion underlying the present approach to Image Understanding is that it should be carried out in the context of a model of the objects of interest in the scene. That model will be created and refined based on a comparison between the given image and a synthetic image of the model. In order that this comparison should yield the most information it is necessary that the synthesized images should be as realistic as possible, especially in those characteristics that are perceptually important. Such characteristics include shape, orientation, lighting,

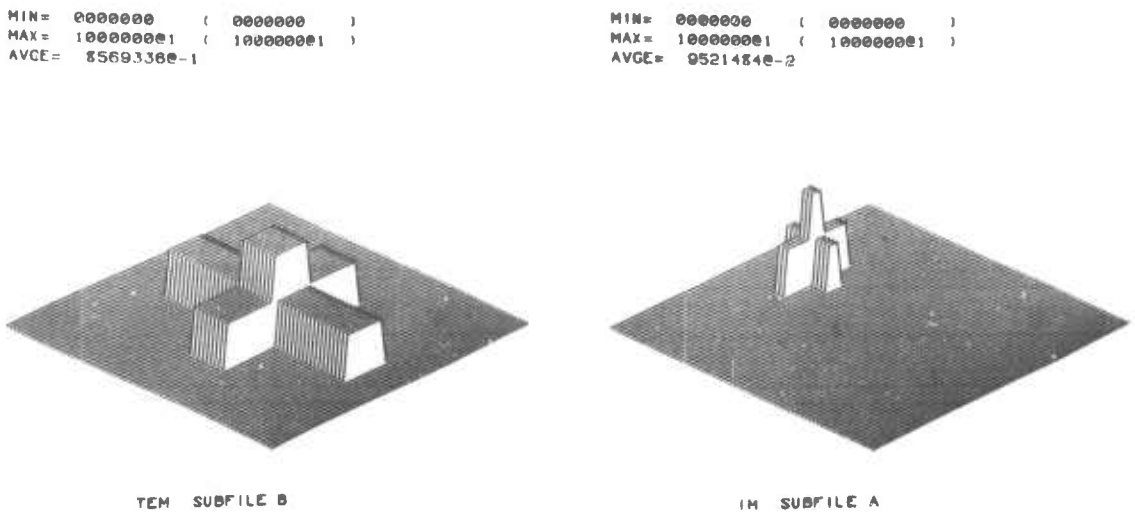

a. Template.

b. Image.

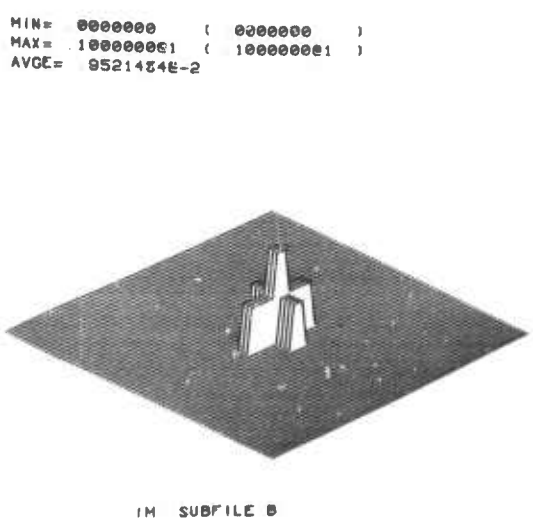

c. Image translated to origin.

Figure 2. Centroid method for generalized 2-D correlation

(continued)

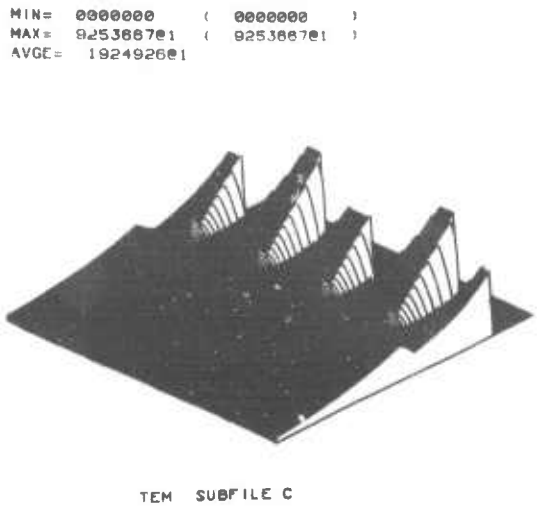

d. Template in exp-rad coords.

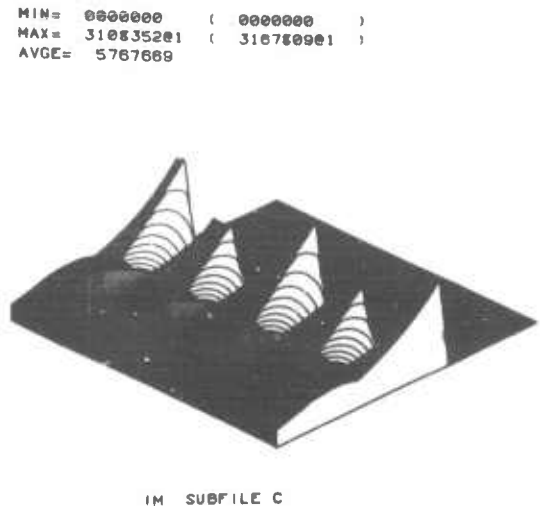

e. Image in exp-rad coords.

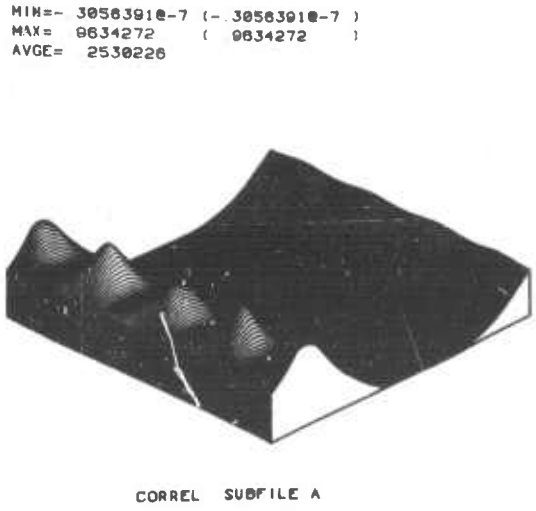

f. Correlation coefficient.

Figure 2. Centroid method for generalized 2-D correlation

highlights, reflections, surface properties, texture color, and shadows.

While current state-of-the-art synthetic images bear a remarkable resemblance to images of real scenes, the simulation of reflections, surface properties, texture and shadows leaves much room for improvement. To this end the development of improved capabilities in these areas has been pursued.

Texture

Most image synthesis techniques model surfaces as being smooth and uniformly colored. Even with the use of highlighting techniques the surfaces appear to be made of some plastic type material. Catmull's work (1974)[5] introduced a technique for mapping patterns onto a surface for the purposes of display. The technique simulates the painting of a pattern onto a surface. No restriction was made on the form of the pattern. Consequently, textured surfaces could be simulated by the use of appropriate patterns.

As in all discrete implementations of continuous signals, an aliasing problem arises in the process of mapping a pattern onto a surface. This is a compound problem in this case since the texture pattern is sampled to determine the shade of a particular piece of surface, and

the surface is sampled for the purposes of display. As with all aliasing problems, the solution is to band limit the signal being sampled by low pass filtering. However, the problem of band limiting the mapped pattern is severe since it is sampled at non-uniform intervals. A paper by Blinn and Newell (1976)[6] presents a new technique for carrying out the pattern mapping process with a minimum of aliasing problems. This new technique is computationally more efficient than previously used methods, though techniques for carrying out this type of process in real time have not yet been developed. This paper is included.

Reflections

The correct simulation of mirror reflections is necessary when dealing with highly polished or glazed objects. Such reflections give rise to highlights, which are perceptually important shape clues. Previous attempts to simulate reflections have been limited to the reflection of point light sources in surfaces which are polished, but not mirror reflecting.

A new technique for simulating true mirror reflections in curved surfaces has been developed. This technique is also described in Blinn (1976). Basically the algorithm functions by recursively fragmenting the surface in the same manner as is used in the patterm mapping algorithm. When a

fragment is to be displayed, its surface normal is used to address a cylindrical projection of the environment surrounding the object. The intensity value so obtained is used to scale the white specular reflection component of the surface shading. The diffuse reflection component is computed using one of several previously reported shading rules, such as Lambert or Bui-Tuong (1975)[7].

2.4. Advanced Modelling systems

The development of an analysis by synthesis system for Image Understanding relies heavily on the existence of a three-dimensional model of the scene being analyzed. This model will be updated and refined as the analysis proceeds.

In order to reduce the number of degrees of freedom in such a model, the use of a single-primitive modelling system is untenable. For example, a commonly used primitive for modelling three-dimensional scenes is the planar polygon. However, the modelling of even comparatively simple scenes involves the use of many hundreds, or even thousands, of polygons. Clearly, higher level, parameterized, models are needed.

The use of higher level parameterized models implies specialization of the modelling system for each type of model. This is because the types of parameters, and the type of control that they give over the model, will vary

widely from one model to another. However, such specialization is undesirable in that many different modelling systems will be required, and the modelling of a scene containing more than one type of model becomes impossible.

Consequently, system structures have been developed which permit several high level parameterized mcdels to coexist within the same system, and to communicate with one another.

The system structure that has been developed and implemented is based on ideas from several sources, including Winograd (1973)[8], Hewitt (1973)[9], SIMULA (1973)[10], Newell (1975)[11]. The essential idea is that objects should be represented as combinations of data and procedure, where both the data and procedure parts can vary from object to object. This is in contrast to the more conventional use of only data to represent objects, all objects being interpreted by a common centralized set of procedures.

The added flexibility afforded by this approach is just what is needed in a geometric modelling system for use in Image Understanding. However, the implementation of such a system raises many questions, such as: how should these models be created and manipulated, at which level should such models communicate, what facilities in the system are

independent of the models themselves? In an attempt to answer these questions a prototype system has been implemented based on the system structure shown in Figure 3.

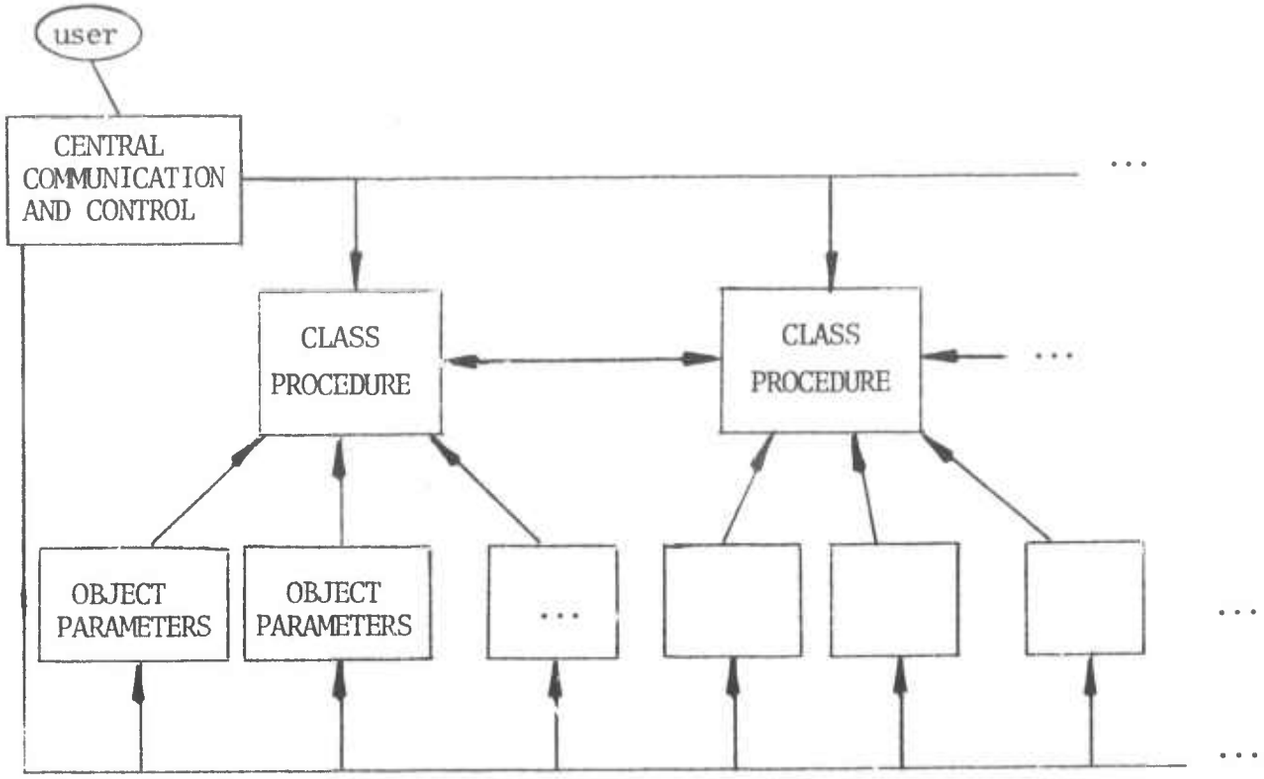

Figure 3.

While each object is conceptually represented as a procedure and data, the fact that several objects may be of the same type leads us to use instances of a class of object. The class is identified by the procedure part which is shared by the objects of the class. The data parts of the objects in a class are held separately. Each block of

object data can be thought of as a set of parameters to the class procedure, and, together with the class procedure it defines the object. The combination of class procedure plus object parameters can be thought of as a single-primitive subsystem residing within the total system.

Since the class procedure is shared among the objects of the class, the only unique manifestations of the objects are their data blocks. Consequently, the central communication and control procedures access objects via their data blocks, which in turn, refer to the relevant class procedure. However, for the purposes of creating a new instance of a class, the central procedures must communicate directly with the class procedure.

The nature of the messages that are exchanged between class procedures and the central control procedures needs careful attention. In a conventional, single-primitive system, the communication between the central procedures and the data base is in terms of the primitives of the system. However, in the structure used here the communication must be in terms of much higher level abstractions. For example, if it is necessary to synthesize an image of an object for the purposes of comparison with a given photograph, the message "give me your image" would be sent to the object (class procedure and data). It would then be the task of the class procedure to create the image, possibly using some

commonly available image synthesis utility procedure, and transmit it back to the central procedures. The point is that the technique used to create the image is completely within the control of the class procedure, and can exploit any special information known about the class of objects represented. Examples of such information are symmetry, convexity, and separability.

Another form of communication takes place in the system when one class procedure needs to access another. Many objects in the real world are structured assemblies of other objects. In such cases the class procedure would make reference to the sub-objects by the same mechanism as that used by the central procedures. For example, a request for the weight of such an object would result in a depth first traverse of the tree of sub-objects.

The creation and modification of objects is carried out by the relevant class procedure. The creation and modification of class procedures is not so straightforward. Not many programming systems provide for dynamic modification of program. Those that do still leave much to be desired at the human engineering level. Consequently, the preliminary design of a new programming language has been completed. Programs written in this language would be run interpretively, to achieve the dynamic capabilities. Moreover, the primitives of this language include the

abstractions dealt with in the system. These include objects, class procedures, instance blocks, tree structures of objects, etc. in addition to the conventional arithmetic and control capabilities. This work is the subject of an ongoing study.

3. Further Work

The purpose of this year's work was to develop the tools necessary in an automated or semi-automated image understanding system. While significant progress was made in the four areas described, further work of this type is still needed. In the area of abstracting perceptually relevant information, the properties and application of the Mandala transform need to be investigated.

In the area of generalized correlation the techniques developed need extending to handle multiple objects and objects on textured backgrounds. Extensions to the 3-D problem are also needed. In this, two extra rotational degrees of freedom are introduced. This invokes the need to use a 3-D template which is a model of the object in question. For any trial orientation of the model a synthetic image would be generated and correlated with the given image using the methods described here.

This brings up the need for high fidelity synthetic images. Again, while significant advances have been made,

especially in the area of realistic surface properites, further work is needed. However, the quality of images now being produced is good enough for preliminary use in an integrated image understanding system. The main outstanding deficiency is still the simulation of shadows, which give important perceptual clues and should therefore be included.

The area of modelling systems needs little further development. The system that was implemented was a trial testbed and as such would not be suitable for a working integrated system.

The next step in this work would be to integrate the various parts in an image understanding system. The modelling system would store and manipulate 3-D models of objects of interest. These would be used to synthesize 2-D images for comparison with the given image. Image enhancement techniques would be applied to both the given and synthesized images to concentrate attention on perceptually important features. Generalized correlation techniques could then be used to find the best fit between real and synthetic images, and to measure the quality of the fit. Feedback from the correlation, either manual or automatic, would be used to modify the model and complete the iterative cycle.

communications

acm

In this Issue:

Papers from Conference on Computer Graphics

Graphics and
Image Processing

Texture and Reflection in Computer Generated Images

James F. Blinn and Martin E. Newell
University of Utah

In 1974 Catmull developed a new algorithm for rendering images of bivariate surface patches. This paper describes extensions of this algorithm in the areas of texture simulation and lighting models. The parametrization of a patch defines a coordinate system which is used as a key for mapping patterns onto the surface. The intensity of the pattern at each picture element is computed as a weighted average of regions of the pattern definition function. The shape and size of this weighting function are chosen using digital signal processing theory. The patch rendering algorithm allows accurate computation of the surface normal to the patch at each picture element, permitting the simulation of mirror reflections. The amount of light coming from a given direction is modeled in a similar manner to the texture mapping and then added to the intensity obtained from the texture mapping. Several examples of images synthesized using these new techniques are included.

Key Words and Phrases: computer graphics, graphic display, shading, hidden surface removal

CR Categories: 3.41, 5.12, 5.15, 8.2

A version of this paper was presented at SIGGRAPH '76: The Third Annual Conference on Computer Graphics, Interactive Techniques, and Image Processing, The Wharton School, University of Pennsylvania, July 14–16, 1976.

This work was supported in part by ARPA under Contract DAH15-73-C-0363. Author's address: Computer Science Department, University of Utah, Salt Lake City, UT 84112.

Introduction

In 1974 Edwin Catmull [2] developed an algorithm for rendering continuous tone images of objects modeled with bivariate parametric surface patches. Unlike most earlier algorithms [6, 8, 9, 10], which require that objects be approximated by collections of planar polygons, Catmull's algorithm works directly from the mathematical definition of the surface patches. The algorithm functions by recursively subdividing each patch into smaller patches until the image of each fragment covers only one picture element. At this stage, visibility and intensity calculations are performed for that picture element. Since the subdivision process will generate picture elements in a somewhat scattered fashion, the image must be built in a memory called a depth buffer or Z-buffer. This is a large, random access memory which, for each picture element, stores the intensity of the image and the depth of the surface visible at that element. As each patch fragment is generated, its depth is compared with that of the fragment currently occupying the relevant picture element. If greater, the new fragment is ignored, otherwise the picture element is updated.

This paper describes extensions of Catmull's algorithm in the areas of texture and reflection. The developments make use of digital signal processing theory and curved surface mathematics to improve image quality.

Texture Mapping

Catmull recognized the capability of his algorithm for simulating variously textured surfaces. Since the bivariate patch used is a mapping of the unit square in the parameter space, the coordinates of the square can be used as a curvilinear coordinate system for the patch. It is a simple matter for the subdivision process to keep track of the parameter limits of each patch fragment, thereby yielding the parameter values at each picture element. These parameter values may then be used as a key for mapping patterns onto the surface. As each picture element is generated, the parametric values of the patch within that picture element are used as input to a pattern definition function. The value of this function then scales the intensity of that picture element. By suitably defining the pattern function, various surface textures can be simulated.

As Catmull pointed out, simply sampling the texture pattern at the center of each picture element is not sufficient to generate the desired picture, since two adjacent picture elements in the image can correspond to two widely separated points in the patch parameter space, and hence to widely separated locations in the texture pattern. Intermediate regions, which should somehow influence the intensity pattern, would be skipped over entirely. This is a special case of a phe-

nomenon known as "aliasing" in the theory of digital signal processing. This theory [7] treats the image as a continuous signal which is sampled at intervals corresponding to the distance between picture elements. The well-known "sampling theorem" states that the sampled picture cannot represent spatial frequencies greater than 1 cycle/2 picture elements. "Aliasing" refers to the result of sampling a signal containing frequencies higher than this limit. The high spatial frequencies (as occur in fine detail or sharp edges) reappear under the alias of low spatial frequencies. This problem is most familiar as staircase edges or "jaggies." In the process of texture mapping, the aliasing can be extreme, owing to the potentially low sampling rate across the texture pattern.

To alleviate this problem we must filter out the high spatial frequency components of the image (in this case the texture pattern) before sampling. This filtering has the effect of applying a controlled blur to the pattern. This can be implemented by taking a weighted average of values in the pattern immediately surrounding the sampled point. Digital image processing theory provides a quantitative measure of the effectiveness of such weighting functions in terms of how well they attenuate high frequencies and leave low frequencies intact.

Catmull achieved the effect of filtering by maintaining an additional floating-point word for each picture element. This word contained the fraction of the picture element covered by patch fragments. For each new fragment added to the picture element, the texture pattern was sampled and the intensity was averaged proportionally to the amount of the picture element covered by the patch fragment. Examination of the spatial frequency filter effectively implemented by this technique shows that it is much better than point sampling but is not optimal.

The method discussed here does not require the extra storage and uses a better anti-aliasing filter. This filter is implemented by a weighting function originally used by Crow [3] to minimize aliasing at polygon edges ("jaggies"). It takes the form of a square pyramid with a base width of 2×2 picture elements. In the texture mapping case, the 2×2 region surrounding the given picture element is inverse mapped to the corresponding quadrilateral in the u,v parameter space (which is the same as the texture pattern space), see Figure 1. The values in the texture pattern within the quadrilateral are weighted by a pyramid distorted to fit the quadrilateral and summed.

Fig. 1. Region of texture pattern corresponding to picture element: left-hand side shows texture; right-hand side shows image.

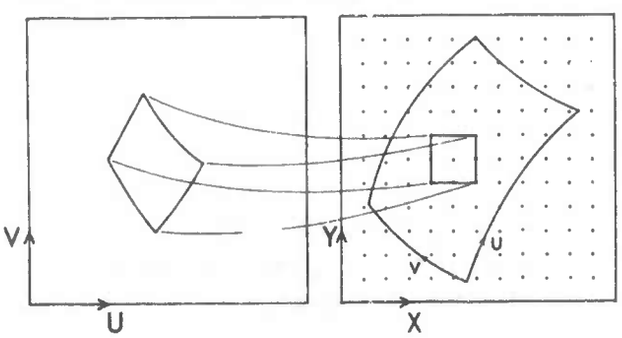

The derivation of the quadrilateral on the texture pattern makes use of an approximation that the parametric lines within one picture element are linear and equally spaced. The X,Y position within a picture element can then be related to the u,v parameters on the patch by a simple affine transformation. This transformation is constructed from the u,v and X,Y values which are known exactly at the four corners of the patch fragment.

Given this algorithm, we now investigate the effects of various texture definition methods. We will use, as our sample object, a plain teapot constructed of 26 bicubic patches. First, the pattern may be some simple function of the u,v parametric values. A useful example of this is a simple gridwork of lines. The result is as though parametric lines of the component patches are painted on the surface, Figure 2. Note that the edges of the pattern lines show very little evidence of aliasing in the form of staircases. Second, the pattern may come from a digitized hand drawn picture, Figure 3. Third, the pattern may come from a scanned-in photograph of a real scene, as in Figure 4. Incidentally, this picture makes the individual patches very clear. This type of pattern definition enables the computer production of "anamorphic" pictures. These are pictures which are distorted in such a way that when viewed in a curved mirror the original picture is regenerated, Figure 5. The patch itself is defined so that the parametric lines are stretched in approximately the correct fashion and a real photograph is mapped onto the patch. Figure 5 should be viewed in a cylindrical mirror (e.g. a metal pen cap) with the axis perpendicular to the page. The fourth source of texture patterns shown here is Fourier synthesis. A two-dimensional frequency spectrum is specified and the inverse Fourier transform generates the texture pattern. This is a simple way of generating wavy or bumpy patterns. Certain restrictions on the form of the input spectrum must be followed to ensure that the pattern has an even distribution of intensities and is continuous across the boundaries. An example of this type of texture is shown in Figure 6. The texture patterns used here were generated before picture synthesis began and stored as 256×256 element pictures in an array in random access memory.

Reflection in Curved Surfaces

The second topic discussed in this paper concerns lighting models. Typically, visible surface algorithms

Fig. 2. Simple gridwork texture pattern: left-hand side shows texture pattern; right-hand side shows textured object.

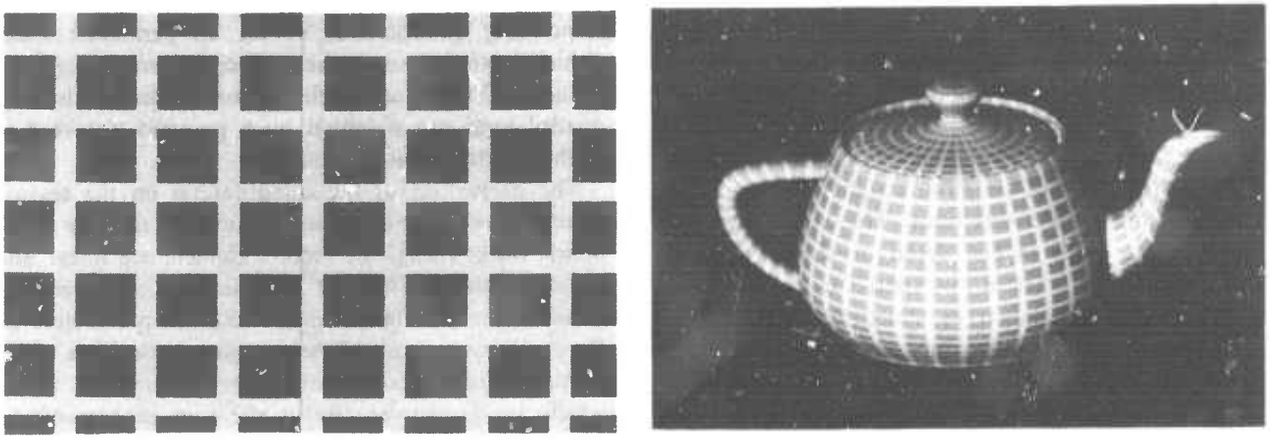

Fig. 3. Hand sketched texture pattern: left-hand side shows texture pattern; right-hand side shows textured object.

Fig. 4. Photographic texture pattern: left-hand side shows texture pattern; right-hand side shows textured object.

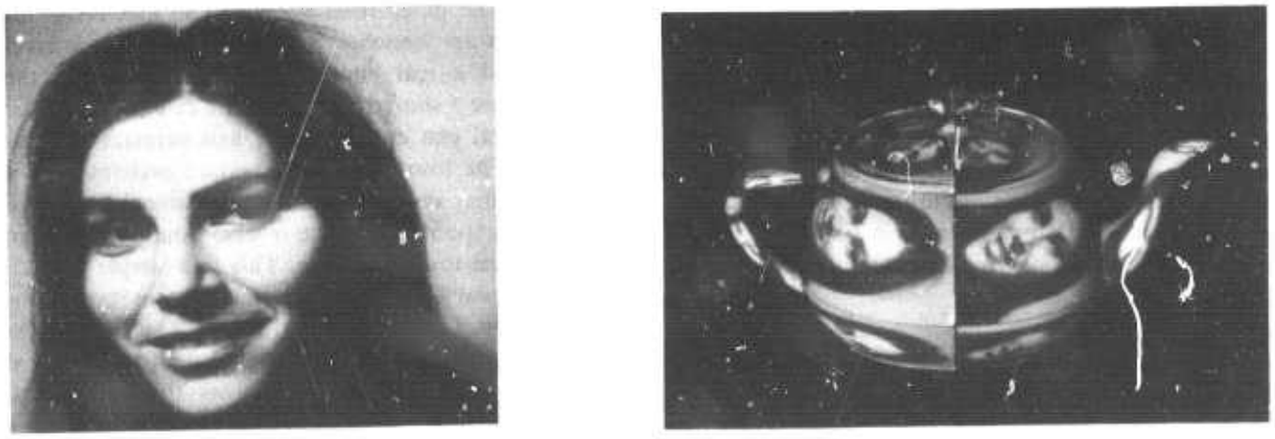

determine intensities within an image by using Lambert's (cosine) law: $i = s(L \cdot N)$, where i = intensity, s = surface shade, L = light direction vector, N = surface normal vector, and "$\cdot$" denotes vector inner product.

Variants on this function, such as

$i = s(L \cdot N)**n$, for $n > 1$

have been used to give the impression of shiny surfaces, but there is little physical justification for such functions, and the range of effects is limited. The modeling of more realistic lighting was first investigated by Bui-Tuong Phong [1]. His model of reflection incorporated a term which produced a highlight over portions of the surface where the normal falls midway between the light

source direction and the viewing direction. This is motivated by the fact that real surfaces tend to reflect more light in a direction which forms equal angles of incidence and reflectance with the surface normal. This can be easily implemented by simulating a virtual light source in a direction halfway between the light source and viewing directions, and raising the result to some high power to make the highlights more distinct:

$$i = s(L \cdot N) + g(L' \cdot N)^{**}n$$

where L' = virtual light source direction, g = glossiness of surface (0 to 1). Figure 7 shows an image generated using the above function with $n = 60$. These techniques work well for satin type surfaces but images of highly polished surfaces still lack realism. This is largely due to the absence of true reflections of surrounding objects and distributed light sources.

The simulation of reflections in curved surfaces requires an accurate model of the properties of the surface and access to accurate normal vectors at all points on the surface. The approximation of curved surfaces by collections of planar polygons is inadequate for this purpose, so extensions of the techniques of Gouraud [5] and Bui-Tuong Phong [1] hold little promise.

The subdivision algorithm, however, provides accurate information about surface position and can be made to give accurate surface normals at every picture element. This is the first algorithm that provides the appropriate information for the simulation of mirror reflections from curved surfaces. For each picture element, the vector from the object to the observer and the normal vector to the surface are combined to determine what part of the environment is reflected in that surface neighborhood. It can be shown that, for surface normal vector (Xn, Yn, Zn) and viewing position (1, 0, 0), the direction reflected, (Xr, Yr, Zr), is

$$Xr = 2*Xn*Zn, \quad Yr = 2*Yn*Zn, \quad Zr = 2*Zn*Zn - 1,$$

Having established the direction of the ray which is reflected to the eye, it remains to find what part of the environment generated that ray. For this, a model of the environment is needed which represents surrounding objects and light sources. Clearly, the view of the environment as seen from different points on the surface will vary. However, if it is assumed that the environment is composed of objects and light sources which are greatly distant from the object being drawn, and that occlusions of the environment by parts of the object itself are ignored, then the environment can be modeled as a two-dimensional projection surrounding the drawn object. Stated another way, the object is positioned at the center of a large sphere on the inside of which a picture of the environment has been painted. These simplifications allow the environment to be modeled as a two-dimensional intensity map indexed by the polar coordinate angles of the ray reflected

Fig. 5. Anamorphic image.

Fig. 6. Fourier synthesis of texture: top shows texture pattern; bottom, texture object.

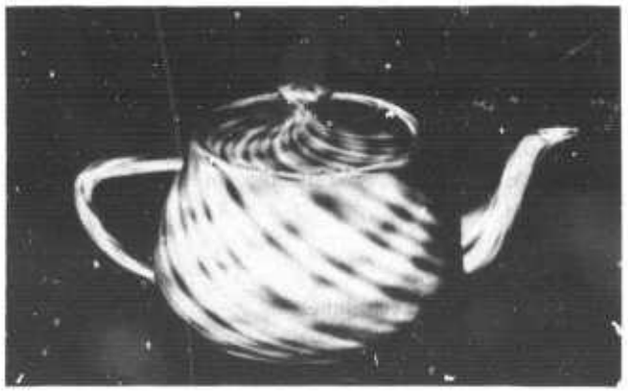

Fig. 7. Plain teapot with highlights.

Fig. 8. Computer generated reflections: left-hand side shows environment map; right-hand side shows reflection in teapot.

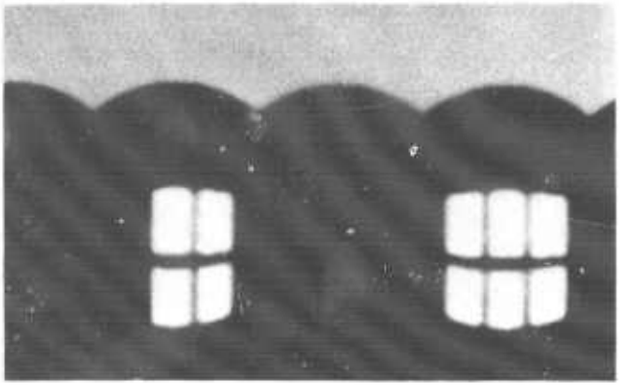

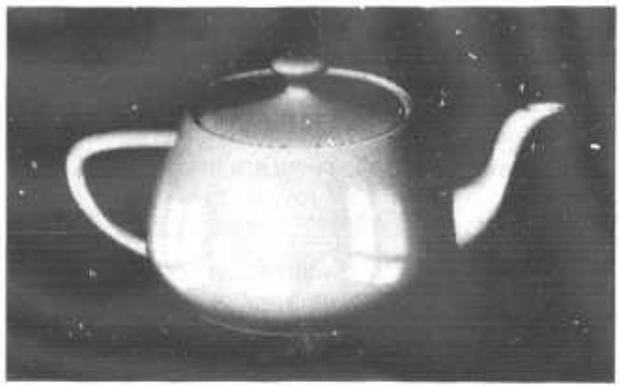

(*Xr*, *Yr*, *Zr*). (Such maps will be shown with azimuthal angle plotted as abscissa and polar angle plotted as ordinate.)

When a reflection direction is computed, it is converted to polar coordinates and the reflected light intensity for that direction is read from the map. This is similar to the technique used to map texture onto a surface, except that the reflection direction, instead of parametric surface position, determines the coordinates in the map. Figure 8 shows an image generated using these techniques.

Use of the surface normal alone is tantamount to modeling the environment on an infinitely large sphere. This has the undesirable effect that the reflected intensity at all silhouette points on the object is the same, and corresponds to the intensity of the environment model diametrically opposite the eye. This deficiency can be corrected by using both the surface normal and surface position to determine what part of the environment map is reflected in a given surface fragment.

Combinations of Techniques

The techniques described above for simulating texture and reflection can be combined to produce images of objects having patterned shiny surfaces. When highlighting is combined with texture mapping, only the component from the real light source should be scaled. This models the highlight as being specularly reflected at the surface and not being affected by the pigment within it. In Figure 9 note how the highlights wash out the texture pattern underneath them. The technique for texture mapping actually keys the texture to the surface so that it moves with the object. Some other techniques, which essentially apply texture to the 2-D image, do not have this property. Note, in Figure 9, how the highlights hardly move with the teapot, whereas the texture does.

Given the texture mapping technique and the environment reflecting technique, we can combine them

Fig. 9. Textured object with highlights, two orientations.

Fig. 10. Highly glazed patterned teapot.

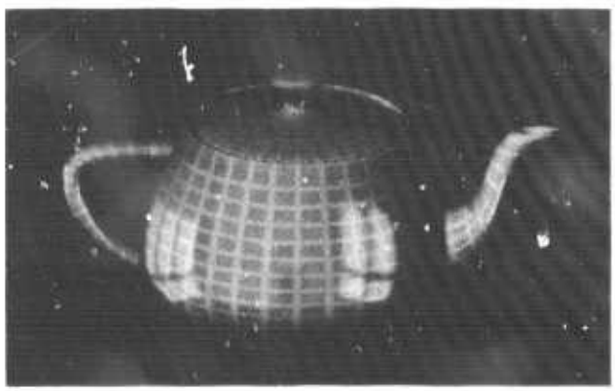

to produce an image of a highly glazed patterned teapot, as in Figure 10.

Resource Requirements

The images shown in this paper were all generated on a PDP-11/45 computer having a 256K-byte random access frame buffer which was used as the depth buffer. The main routines were written in Fortran and the critical parts were written in assembly language. The computation time of the extended subdivision algorithm is roughly proportional to the area covered by visible objects. Images of nontextured objects of the type used in this paper take about 25 minutes. The addition of texture or reflection increases this time by about 10 percent. All images have a resolution of 512×512 picture elements.

Conclusions

By refining and extending Catmull's subdivision algorithm, images can be generated having a far higher degree of naturalness than was previously possible. These generalizations result in improved techniques for generating patterns and texture, and in the new capability for simulating reflections.

References

1. Bui-Tuong Phong. Illumination for computer generated images. *Comm. ACM 18*, 6 (June 1975), 311–317.
2. Catmull, E.A. Computer display of curved surfaces. Proc. Conf. on Comptr. Graphics, Pattern Recognition, and Data Structure, May 1975, pp. 11–17 (IEEE Cat. No. 75CH0981-1C).
3. Crow, F.C. The aliasing problem in computer-synthesized shaded images. Tech. Rep. UTEC-CSC-76-015, Dep. Comptr. Sci., U. of Utah, Salt Lake City, Utah, March 1976.
4. Forrest, A.R. On Coons and other methods for the representation of curved surfaces. *Computer Graphics and Image Processing 1* (1972), 341.
5. Gouraud, H. Computer display of curved surfaces. Tech. Rep. UTEC-CSC-71-113, Dep. Comptr. Sci., U. of Utah, Salt Lake City, Utah, June 1971.
6. Newell, M.E., Newell, R.G., and Sancha, T.L. A solution to the hidden surface problem. Proc. ACM 1972 Ann. Conf., Boston, pp. 443–450.
7. Oppenheim, A.V., and Schafer, R.W. *Digital Signal Processing*. Prentice-Hall, Englewood Cliffs, N.J., 1975, pp. 26–34.
8. Sutherland, I.E., Sproull, R.F., and Schumaker, R.A. A characterization of ten hidden-surface algorithms. *Computing Surveys 6*, 1 (March 1974), 1–55.
9. Warnock, J.E. A hidden-line algorithm for halftone picture representation. Rep. TR 4-15, Dep. Comptr. Sci., U. of Utah, Salt Lake City, Utah, 1969.
10. Watkins, G.S. A real-time visible surface algorithm. Tech. Rep. UTEC-CSC-70-101, Dep. Comptr. Sci., U. of Utah, Salt Lake City, Utah, June 1970.

REFERENCES

[1] Hueckel, M.H. "An Operator Which Locates Edges in Digitized Pictures", JACM, Vol 18, 1, January 1971.

[2] Griffith. "Edge Detection in Simple Scenes", IEEE Trans. on Comp., April 1973.

[3] Robbins, G.M. and Huang T.S. "Inverse Filtering for Linear Shift-Variant Imaging Systems", Proc. IEEE, Vol. 60, 7, July 1972.

[4] Kajiya, J. "Proc. of Symposium on Current Mathematical Problems in Image Science", Monteray Naval Post Graduate School, ONR Sponsorship, Nov. 1976.

[5] Catmull E. "A Subdivision Algorithm for Computer Display of Curved Surfaces", Technical Report UTEC-CSc-74-133, Univ. of Utah, 1974.

[6] Blinn J.F. and Newell M.E. "Texture and Reflection in Computer Generated Images", CACM, Vol. 19, 10, October 1976.

[7] Bui-Tuong, P. "Illumination for Computer Generated Images", Comm. ACM. Vol. 18, 6, June 1975.

[8] Winograd, T. Understanding Natural Language. New York, Academic Press, 1973.

[9] Hewitt C., Bishop P., and Steiger R., "A Universal Modular ACTOR Formalism for Artificial Intelligence", Proc. Third Int. Joint Conf. on Artificial Intelligence, August 1973.

[10] Birtwistle G.M., Dahl O.J., Myhrhaug B., and Nygaard U. SIMULA BEGIN Philadelphia, Auerbach Publishers Inc., 1973.

[11] Newell M.E. "The Utilization of Procedure Models in Digital Image Synthesis". Technical Report UTEC-CSc-76-218, Univ. of Utah, 1975.

SECTION IV

AUDIO PROCESSING

At the end of September 1976 the audio processing work being pursued under this contract became separately funded under ARPA order 3301, contract N00173-77-C-0041 with Naval Research Labs. The small amount of prepatory work done during the three months July 1976 - September 1976 is reported as part of the first semi-annual technical report under that contract: Report Number UTEC-CSc-77-090.

PUBLICATIONS AND PRESENTATIONS

[1] Kajiya, J. "Group Representations and the Modeling of Human Visual Perception" ONR Conference on Mathematics and Image Science, Monterey, Calif. October 1976.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: UTEC-CSc-77-118
2. GOVT ACCESSION NO.:
3. RECIPIENT'S CATALOG NUMBER:
4. TITLE *(and Subtitle)*: SENSORY INFORMATION PROCESSING.
5. TYPE OF REPORT & PERIOD COVERED: Final Technical Report, 1 Jul [redacted]76-31 Mar [redacted]77
6. PERFORMING ORG. REPORT NUMBER:
7. AUTHOR(s): Thomas B. Stockham, Jr.
8. CONTRACT OR GRANT NUMBER(s): DAHC15-73-C-Ø363
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Computer Science Department, University of Utah, Salt Lake City, Utah 84112
10. [redacted] ARPA Order [redacted]-2477
11. CONTROLLING OFFICE NAME AND ADDRESS: Defense Advanced Research Projects Agency, 1400 Wilson Blvd., Arlington, Virginia 22209
12. REPORT DATE: Apr[redacted]77
13. NUMBER OF PAGES: ~~104~~ 116p.
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)*

This document has been approved for public release and sale; its distribution is unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

Same

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

atmospheric turbulence blur, human sensory information processing, image understanding system, harmonic analysis, translation, dilatation, Mandala transform, long-time averaging, interferometric methods, generalized correlation

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

This final report covers the nine month period from July 1976 through the termination of the contract at the end of March 1977. As such, this report describes the accomplishments of this period in detail and thus includes the semi-annual report of the period July 1976 - December 1976. of This report

Section I discusses the theoretical underpinnings of a new mathematical model describing some of the perceptual characteristics of human sensory (cont)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

404 949

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

20. Abstract (cont)
information processing. While much of the mathematics of the proposed model is explained here, experimental testing and verification has yet to take place. This work will continue to be pursued under alternate funding.

Section II discusses in detail the background considerations necessary to understand the problems of removing atmospheric turbulence blur from images. This area of research is currently in transition from the stage of artificial computer models, to the stage of working with the simplest real data, i.e. case of a point light source, a star, viewed through the earth's atmosphere. After the termination of this contract this work will be funded by the Air Force Office of Scientific Research under grant number 77-3212.

Section III reports the background and progress to date on building the tools necessary for an image understanding system. These tools as described were incorporated into a manual image understanding system designed to do analysis by synthesis. The system handles two-dimensional single objects on non-textured backgrounds, but only at the level of a trial testbed. Necessary next steps in developing a working system are outlined.