

THE APL IMAGE PROCESSING LABORATORY

A central Image Processing Laboratory provides APL with a powerful resource for the advancement of ongoing programs in missile technology, space science, oceanography, and biomedical image analysis. Its present and proposed capabilities are discussed, and several applications are described.

INTRODUCTION

The term *image processing* means different things to different people. As used here, it means the computer-aided process of extracting and presenting information from imagery. The process may be a simple matter of improving the appearance of the image. It may also be a more sophisticated procedure for extracting information about image features.

Images come to the Image Processing Laboratory in different forms and from a variety of sources. Many image processing tasks require that the image first be translated into a form that can be accepted as input to a digital computer for analysis. Program objectives frequently require that an enhanced or transformed image be displayed via computer output for viewing and analysis.

AREAS OF IMAGE PROCESSING

It is convenient to divide the capabilities of the APL Image Processing Laboratory (Fig. 1) into three basic areas of image processing:

Image Digitizing. Images are subdivided into an orthogonal array of square picture elements (pixels)

that may be as small as 12.5 micrometers on a side. In the case of monochrome imagery, the optical density of each pixel is measured and recorded as an 8-bit byte of data that defines the density as one of 256 levels. In the case of color imagery, the process is performed three times, using red, green, and blue filters in sequence to create a separate file of data for each primary color.

Digital Image Processing. Digitized imagery, either monochrome or in color, may be processed further using microcomputers, minicomputers, or the mainframe computers in APL's McClure Computing Center. Extensive software exists to support many tasks, and special application software is written as needed.

Digital Image Output. Results may be presented as a computer-generated listing or graph, or as a computer-generated image. Digital imagery, either monochrome or in color, may be displayed on cathode ray tubes or recorded photographically as prints, transparencies, movies, and microfiche.

FUNCTIONAL CAPABILITIES

With the hardware and software capabilities described briefly above, a wide variety of functional capabilities become available. Some of these are:



Figure 1 — A partial view of the APL Image Processing Laboratory showing the Comtal system (left foreground), Matrix camera (left rear), Grinnell workstation (center background), PDP-11/70 computer (right foreground), and Dicomed system (right rear).

Enhancement. Images may be enhanced by improving contrast and sharpening edges (Fig. 2). Artifacts can be removed and graphical overlays created. Since the image digitizer can recognize many more gray levels than the human eye, it is sometimes possible with computer look-up tables to translate otherwise invisible images into a range of densities that can be seen.

Pseudocolor. The recognition of features in a black and white image can frequently be improved by assigning colors to different gray levels in the image and producing a new "pseudocolor-encoded" image. Interactive image processing equipment allows an analyst to vary the color code in real time and observe the results on a display screen in order to bring out the desired features.

Convolution. Sometimes it is desirable to filter or smooth an image by the process of convolution, wherein each pixel in the image is modified by a weighted average of the pixels surrounding it. In its simplest form, this process employs a kernel consisting of a 3×3 array of pixels, each of which has an assigned weight. To convolve one pixel in an image, the kernel is centered on that pixel and its value is replaced by the weighted average of the values of all pixels in the array. A hardware convolver performs this process on a 512×512 pixel image in one-thirtieth of a second. Larger kernels on the order of 15×15 pixels can be used with software routines controlling the hardware.

Computer Output Microfilm. When large volumes of graphic output are generated, it is often convenient to record the output on film. This can be done on 16 and 35 millimeter roll film in either color or black and white. Even greater reduction can be achieved by recording monochrome data on 105 millimeter microfiche in a variety of formats, allowing up to 420 frames per fiche.

Presentation Graphics. High quality presentation graphics can be generated in black and white or color as 35 millimeter slides, 8×10 inch viewgraphs, or paper prints (Fig. 3). The design of such graphics can be performed using any of the many graphics terminals connected to the IBM 3033 MP computer in the McClure Computing Center.

Animations. Computer-generated animations can be created as loop movies with an interactive processor and viewed at a controlled frame rate on a display monitor. Video tapes may be made from these loop movies. High quality movies can be generated on the computer-controlled film recorder in 16 or 35 millimeter format. Such movies may be of any length, and sound may be added. Figure 4 shows one frame of a movie generated this way.

Transforms. Image analysis frequently requires the spatial transformation of the original image. This can be accomplished readily on a computer once the image has been digitized. It has been found in a study of ocean waves that by taking the Fourier transform of an ocean image one can deduce information on the direction, frequency, and speed of the waves.




Figure 2 — This computer-assisted tomograph (CAT) scan was generated in the Department of Radiology of the Johns Hopkins Hospital for an assessment of the fibrosis that has invaded the liver. Also apparent in this cross-sectional view are the kidneys, spleen, abdominal ribs, and a spinal vertebra. (top) An Optronics filmwriter has recorded an 8-bit representation of the 16 bits of digital data acquired with the American Science and Engineering CAT scanner. (bottom) A contrast enhancement of liver tissue density was computed by scaling the square root of the CAT scan data.

Geometric Correction. Some imagery requires geometric correction to remove distortion introduced by the attitude or transfer characteristics of the sensor. Such corrections can be made using available software. Interactive systems allow the controlled superposition of two images of the same scene taken from different sensors. Corresponding features of the two images can be made to coincide using geometric correction algorithms.

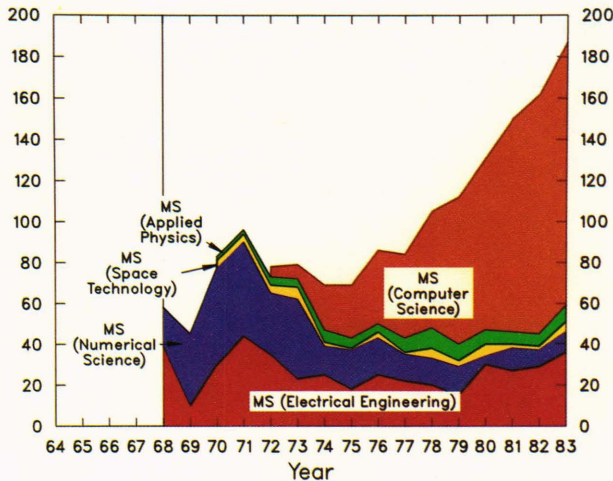


Figure 3 — An example of "data representation" graphics created on 35 millimeter slide film from software on the mainframe computer. This graph shows the number and distribution of Master's degrees awarded by the APL Education Center.




Figure 4 — A 20-minute animated film, *Anomalous Propagation of Radar*, was generated as a tutorial on the physics of radar and the propagation of radar under anomalous atmospheric conditions. The frame shown is from a sequence that illustrates the changes in radar propagation over a 24-hour period. The warming of the atmosphere has caused an elevated duct to intersect the radar site, resulting in a void in radar coverage. This film was prepared under the direction of H. W. Ko.¹ The programming was done by J. P. Skura and C. A. Waters using DISSPLA graphics software.

Feature Extraction. Information can be extracted from an image to describe the size, shape, or location of objects or regions in the image. Such regions can be identified or "classified" by specifying a range of pixel values, possibly in more than one color channel or "band." The pixel value histogram of known image areas can be used to train the classifier. The classified image is displayed with each identified region in a different color or gray value. Edges of objects in the image can be extracted by a convolution filter. Features that can be identified visually may be measured at the cursor position on a video display. A graphic digitizer may be used if the image exists as a

slide, movie, or other copy that is not in digital form. Objects or targets with known shape can be located by using the convolution hardware (template matching) to look for a correlation between the target and image pixels.

HARDWARE AND SOFTWARE

Digitizing Equipment

Graphic data may be digitized on the stand-alone Comp-U-Grid system, which includes a digitizing table and a Nova 2/10 computer. Image digitizing is performed on the Optronics P-1700 drum scanner/writer, which accepts monochrome or color imagery, either transparent or opaque. The features of these digitizers are shown in Table 1.

Table 1 — Digitizing equipment.

Feature	Comp-U-Grid	Optronics P-1700
Data form	Graphics	Imagery
Input medium	Paper or roll film	Paper or cut film
Image size	To 42 × 60 inches	To 10 × 10 inches
Resolution	0.001 inch	12.5 to 400 micrometers

Processing Equipment

Four image processing work stations are available. Two of these are connected to the Comtal Vision One/20 interactive image processor, which employs an embedded DEC LSI-11/13 microprocessor. The Comtal system can be used in stand-alone operation or with an associated DEC PDP-11/70 computer. The other two work stations are connected to a Grinnell GMR-270 system that is interfaced directly to the 3033 computer. The characteristics of these image processing systems are shown in Table 2.

Table 2 — Image processing systems.

Feature	Comtal	Grinnell
Computer	Embedded LSI-11 PDP-11/70 host	IBM 3033 host
Control	Keyboard, track ball, and data tablet	APL or TSO terminal and joystick
Display size	512 × 512 pixels	512 × 512 pixels
Memory size	8 images	4 images
Video digitizer	None	One 8 bit

In addition to the image processing systems described above, two general purpose computers are available. The PDP-11/70 computer serves as the controlling element for most of the image processing

equipment. It is linked to the 3033 computer, which serves APL as a whole. Some of the features of these computers are shown in Table 3.

Table 3 — Digital computers.

Feature	PDP-11/70	IBM 3033 Host
Memory	640 kilobytes	16 megabytes
Disk storage	610 megabytes	15.4 gigabytes
Mass storage	None	105 gigabytes
9-track tape drives	3	8
Maximum tape density	6250 bits/inch	6250 bits/inch
Terminals	6	900
Operating system	RSX-11M plus	MVS

Software Packages

The equipment in the Image Processing Laboratory is supported by software packages for image processing and graphics.

Image Processing. Both the Comtal and Grinnell systems are supported by software designed to make efficient use of the hardware. The interactive capabilities of these systems make their use particularly easy. Figure 5 shows an effective use of the Grinnell system. A selection of the available image processing software is shown in Table 4. The general purpose packages support histograms, profiles, and other statistical measures, contrast enhancement, pseudocolor encoding, graphical overlays, convolution, pixel arithmetic, and logical operations.




Figure 5 — The Grinnell Image Processing System attached to the IBM 3033 MP computer was used to generate an approximate solution to a two-dimensional Laplacian differential equation. The evolution of the solution is displayed on a color monitor. For this image, the differential equation is used to describe the flow of heat from a hot boundary (red) to a cool interior (gray). Intermediate temperatures are highlighted by pseudocoloring the image.

Graphics. Graphics capabilities are a part of SAS, APL, Speakeasy, and various systems for simulation and engineering design. There are subroutine libraries for Fortran and PL/I and many user-written and

Table 4 — Image processing software.

Package	Method of Use	Features	Source
Grinnell	Command language (interactive) Function library (PL/I, APL)	General purpose DISSPLA graphics	Developed at APL
Comtal	Command language (interactive)	General purpose Freehand graphics	Purchased
Smips-Vicar	Command language (batch)	General purpose Any size image Annotation files Geometric correction Classification	Public domain
IEEE signal processing	Function library (Fortran)	Spectral analysis	Public domain
SACS	Menus	Convolution Spectral analysis Contrast enhancement	Developed at APL




Figure 7 — Projected view of the lithium hydride/helium interaction surface truncated at the 0.5 electronvolt contour level. Successive contours (expressed in millihartree) differ by a factor of $\sqrt{10}$ in energy. This representation was produced on the CalComp plotter from digital data computed by D. M. Silver² of the Milton S. Eisenhower Research Center.

3. The Optronics C-4300 writes high resolution color film, also in large format.
4. The Dicomed D48C computer-controlled film recorder writes on a variety of film sizes in both black and white and color.

The output capabilities of the four systems are compared in Table 6.

Table 6 — Hard-copy image output capabilities

Output format	Matrix	Optronics	Dicomed
16 mm roll			X
16 mm jackets			X
16 mm movies			X
35 mm roll	X		X
35 mm slides	X	X	X
35 mm movies			X
105 mm microfiche			X
10 x 10 inch cut film		X	
8 x 10 inch cut film	X	X	X
8 x 10 inch Polaroid	X	X	X
SX-70 Polaroid	X		

System Configuration

The Optronics, Comtal, and Dicomed systems described are all interfaced to the PDP-11/70 computer, which is linked to other DEC computers and to the 3033 computer via the APLnet computer network. A Grinnell work station is located next to the

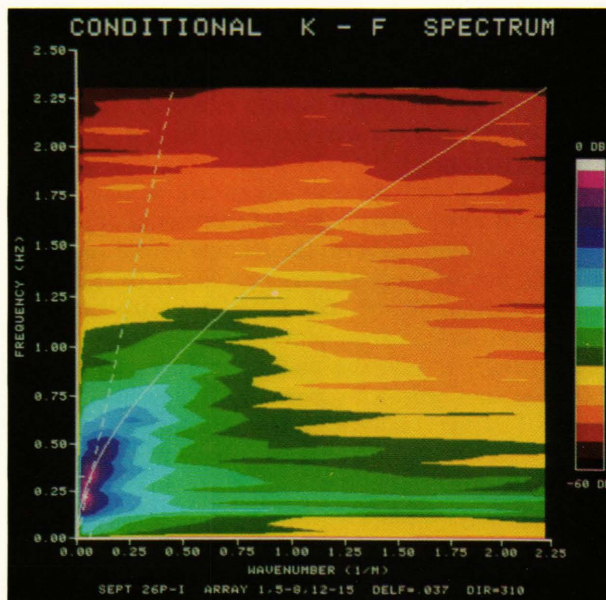


Figure 8 — This two-dimensional wave spectrum depicts the dispersion of surface waves as viewed from an elevated platform with an array of optical sensors. The spectral power as a function of wavenumber and wave frequency has been color coded to assist in comparing experimental data with theoretical models. The data base for this display was created by B. L. Gotwols and was discussed in the *Johns Hopkins APL Technical Digest* 3, 49-58 (1982)

Comtal system, with the two systems sharing the Matrix camera. Digital imagery can be transmitted between computers and peripherals for processing, display, or storage. Many image processing applications have arisen that take advantage of this system configuration. A block diagram of the system configuration is shown in Fig. 9.

APPLICATIONS

Imaging Systems

A number of APL projects involve the development of imaging systems for weapon systems applications. Among these are digital image processing research programs involving Autonomous Target Recognition and Wide Area Guidance and Control. One program is aimed at developing and analyzing optical scene matching algorithms. Such algorithms involve the comparison of small images (the target images) sensed by scanners with large stored images (the reference images). Correlation techniques and other algorithms are then used by missiles to maintain the correct path to the target.

This work required extensive computer analysis of digital image data along with a means of displaying images quickly. In addition, it was considered necessary to use an image processing system with which all users were familiar and to which there was wide access. The analytical and experimental nature of the development also required that access be available through common high-level languages so that results

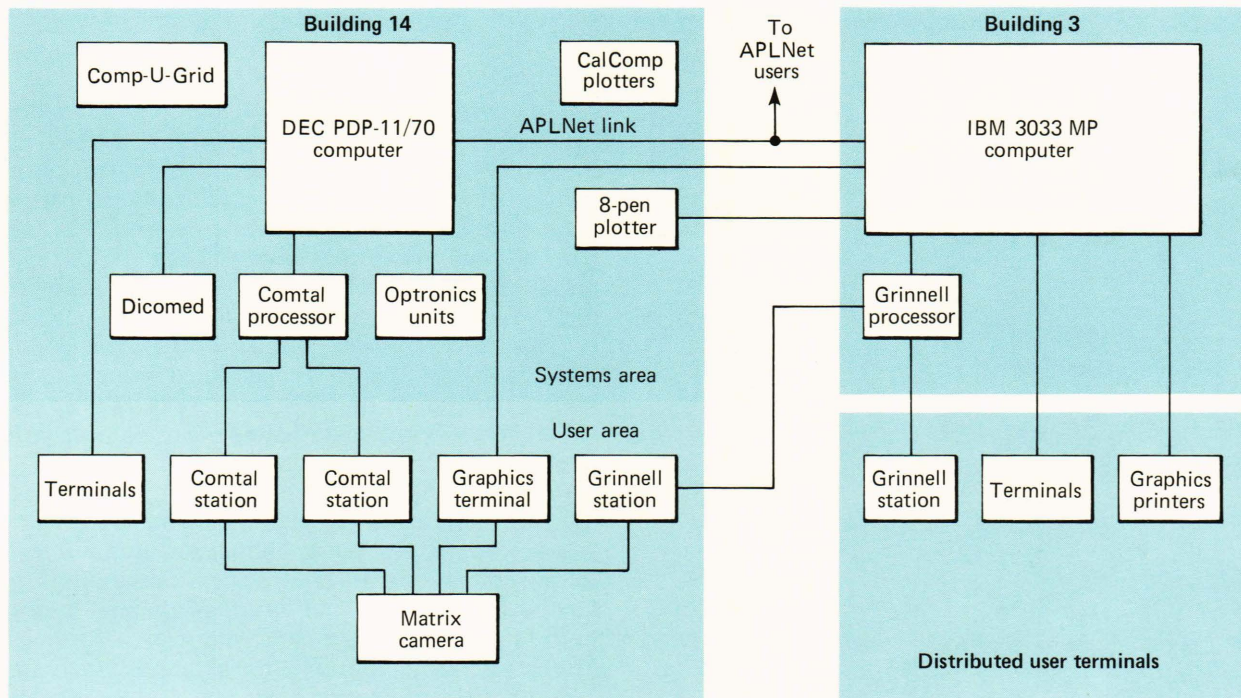


Figure 9 — Block diagram of the APL Image Processing Laboratory. The two main computers, a DEC PDP-11/70 and the IBM 3033, located in separate buildings, are connected by means of the APLNet link, which also connects to other computers. A user area is located in Building 14, and user terminals are distributed throughout APL.

could be achieved quickly without substantial user training. As a result of these requirements, the development has been carried out on the Grinnell Image Processing System interfaced to the 3033 computer.

The primary software system developed on the Grinnell system for this program is the Scene Analysis and Correlation System, which is a completely executable PL/I program that is accessed through TSO on the 3033 computer. It is a menu-driven system, well suited for nonprogrammers who need access to an interactive image processing system. The routines provided by the Scene Analysis and Correlation System include many operations that come under the general categories of analysis and enhancement. These include one- and two-dimensional filtering, contrast modification, statistical analysis, and arithmetic operations. Utilities for reading and writing images and convolution templates are also included.

An important result of this image processing effort is a definition of scene contrast implemented in the APL language. The definition is applied to a sensed image that has been binarized before it is compared with the reference image. This means that the original gray scale image has been thresholded to produce an image that contains only black and white. Under the definition, each pixel in the sensed image is assigned the probability that it has the correct binary value as determined from the reference image stored in the missile. An example of the application of this definition of contrast is given in Fig. 10.

Future work for this program will involve further analysis of the scene-matching progress, including simulation of the operation of the sensors.

Map Geometric Correction

A satellite or aerial photograph of a region can be used to correct errors in a map of the region. The photo provides a location reference for features that can be identified in both images. The method involves selecting a small area around each feature of the map image and looking for a point of high cross-correlation with the photo image. This process may be done manually or by using the convolution hardware in an image processor.

Absolute latitude and longitude scales are usually known for the photo, allowing it to be corrected to the projection of the map before the cross-correlation is performed. However, in an initial trial, the absolute position information was missing, and the map and photo were brought into approximate coincidence by the following method. Both images were digitized on the P-1700 scanner, and an interesting excerpt of 512×512 pixels was selected for study. The map image was superimposed on the photo image using the Comtal display. Then the map overlay was shifted relative to the photo, under control of the data tablet, to achieve coincidence at several control points. The shift amounts were used to determine a least-squares best linear transformation of the map space onto the photo. A program on the 3033 computer was used for the fit and correction. Figure 11





Figure 10 — The Grinnell image processor was used to develop a definition of image contrast in a study of digital scene-matching area correlation. The upper image is a possible reference scene for the correlator. One measure of the contrast in this scene is given by the probability that a missile will sense the correct binary values (black or white), shown in the middle image, as it flies over the reference area. This probability is calculated for each pixel in the reference scene and is displayed as a pseudocolored image in the lower image. White indicates low probability; black indicates high probability.

shows the superposition of the two images before and after the correction was applied.

The Seasat Synthetic Aperture Radar

The Synthetic Aperture Radar (SAR) sensor aboard the Seasat spacecraft demonstrated sensitivity to a variety of oceanic processes, including the development of wind waves, the propagation of internal waves, wave refraction, and wave dispersion by currents and subsurface topography. These periodic disturbances are often characterized by Fourier spectra to provide oceanographers with two-dimensional statistical estimates of wavenumber density. It is important that these spectra be corrected for the point spread function of the SAR imaging process so that researchers will not be misled by the resolution limits of the sensor.

In optics, the impulse response or point spread function describes the spatial blurring of pictures produced by the physical limitations of the image-forming apparatus. Image speckle is the result of the random spatial distribution of scatterers located within a resolution element of the sensor. Although speckle noise is undesirable for the purpose of image interpretation, its existence can be useful for estimating the impulse response of the Seasat imaging system. The Fourier transform of the impulse response function is commonly called the wavenumber response function. Linear theory predicts that the wavenumber response found for a sum of uncorrelated data samples will be the same as that produced by a single impulse. This equivalence is most useful when it is not possible to isolate a point scatterer for measurement of the impulse response at low spatial frequencies. Furthermore, the collective effects of backscattering facets that move at random on the ocean surface cannot be interrogated with an impulse image but can be included in the analysis of speckle noise.

Linear regression of a squared error function has been applied to 3×3 kilometer segments of Seasat images from the desert near Goldstone, Calif. and Chesapeake Bay, yielding similar results. These speckled scenes that have been chosen as characteristic random data have Rayleigh intensity distributions and spatially uniform mean values. Their power spectra, derived with fast Fourier transform computations, have been used to synthesize polynomial functions that approximate the wavenumber response of the SAR system. Figure 12a includes textured regions in Chesapeake Bay that were subjected to Fourier analysis in the Image Processing Laboratory to determine sensor response characteristics from speckle noise.

The random velocities of scatterers on the Bay do not significantly reduce the spectral resolution and bandwidth as measured for stationary scatters in the desert. Furthermore, the wavenumber response of the SAR system seems unchanged over the 32 days separating Seasat orbits 882 and 1339. Generally, it was found that polynomial approximations of the wavenumber response function can be used to re-




Figure 11 — In this example of geometric correction of one image to another, an aerial photograph (top) of Huntsville, Ala. is shown with a superimposed line drawing digitized from a topographic map of the same region. The map was distorted (bottom) to agree with the photo. This was done by a linear mapping of the pixel locations, determined by a least-squares fit to a small number of control points identified in each image.

move the Seasat sensor bias from ocean wave spectra for wavenumbers less than the practical resolution limit of $2\pi/50$ radians per meter.

Raster Movie Display of Seasat Ocean Data

APL's Space Department is engaged in the oceanographic analysis of wind-driven wave systems in the North Atlantic as recorded by the SAR aboard Seasat in 1978. Twenty magnetic tapes were purchased from MacDonald, Dettwiler and Associates via the Can-

ada Centre for Remote Sensing. Each tape contained more than 12 million words of image data representing 1600 square kilometers of ocean surface. The facilities of the Image Processing Laboratory have been employed for data compression and normalization, for spectral analysis of ocean scenes, and for film and video displays of sea surface reflectance and wave spectral density. Image statistics and Fourier transforms have been computed with the PDP-11/70 and the 3033 computers to provide estimates of wind vectors, wavelengths, and directional wave power.³

A 16 millimeter filmstrip was produced with the Dicommed film recorder to display ocean wave spectra, as well as the Seasat ocean scenes (Fig. 13) that were subjected to Fourier analysis. The raster movie begins with a global view of Chesapeake Bay, exercising the film recorder at maximum resolution. The digital scene consists of 3100×4050 pixels representing 1600 square kilometers. Since all of the 4050 pixels in the vertical dimension of the 16 millimeter frame could not be displayed at once, each successive frame was extracted from the digital image with a 50 meter increment along the ground track of the satellite. The margin of the 16 millimeter frame was filled with latitude and longitude annotations along with the satellite direction and North arrows, all written as vector strokes rather than raster pixels.

The movie continues with a scrolled display of a 900×6 kilometer sample of the surface reflectance in the North Atlantic recorded by the SAR during Seasat orbit 1339. The image data were compressed by a factor of four in the writing of magnetic tapes containing Dicommed code for approximately 4700 raster frames, each incrementally shifted 200 meters along the path. Each frame (256×256 pixels) was read from magnetic tape and enlarged by a factor of 10 in both dimensions to fill a 2560×2560 raster unit area on the Dicommed cathode ray tube. The entire monochrome sequence depicts image data acquired in about 2 minutes aboard Seasat.

The scrolled display of the image data is followed by a sequence of ocean wave spectra computed from the same 900 kilometer strip of image data. A total of 147 ocean wave spectra were computed depicting changes in directional wave power and wavelength along the path. Those spectra are pseudocolor encoded for ease of interpretation.

The entire 16 millimeter filmstrip, consisting of approximately 7200 raster frames with vector annotation, can be viewed in 5 minutes with a standard projector running at 24 frames per second. Monochrome frames were written in 28 seconds and 102 seconds on the average when operating the Dicommed film recorder at its lowest and highest raster resolutions, respectively. Color frames were written in 81 seconds at low resolution, requiring three sequential exposures through red, green, and blue filters.

Seasat and Skylab Scenes of Nantucket Island

The classification of terrain near Nantucket Island represents a typical application of the facilities of the



Figure 12a — Chesapeake Bay as observed on September 28, 1978, with the Seasat L-band SAR. The Naval Academy and the city of Annapolis can be distinguished southwest of the twin spans of the Bay Bridge that extends eastward to Kent Island. A tanker south of the Bay Bridge is evidently anchored there since no wake is apparent. Several smaller vessels in the upper portion of the scene do have wakes behind them, and the fastest moving of the group appears to be Doppler shifted off its wake. This scene was processed by the digital correlator at the Jet Propulsion Laboratory and was recorded on film in the APL Image Processing Laboratory using the Optronics P-1700 filmwriter to expose each of 10^7 picture elements.

Image Processing Laboratory. A SAR image of that area (Fig. 14) was recorded by the Seasat satellite during the fall of 1978. A comparison of the Seasat image with an optical photograph of the same region acquired by Skylab (Fig. 15) in the fall of 1976 is interesting for a number of reasons. For example, radar engineers would be interested in the geometric and radiometric characteristics of the Seasat image relative to the Skylab photograph, while oceanographers would compare the scenes to identify similarities in the regional circulation or differences in the bathymetry over a two-year span.

The Seasat image was obtained from the Jet Propulsion Laboratory, where a PDP-11/34 computer and an AP120B array processor were used to con-

struct the digital 5840×6144 pixel image from the time-based radar data. This image was then reduced to 2920×3072 pixels (50 meter resolution) by averaging a 4 pixel area to create a single pixel. A geometric correction was applied with the PDP-11/70 computer to map the slant range coordinate to ground range.⁴ After rotating the Seasat image by 90° , a second geometric correction was applied to remove the azimuth skew.⁵

The Skylab scene of Nantucket Island was acquired as a 10×14 inch color print from archives of the National Oceanographic and Atmospheric Administration. The image was photographically reduced to an 8×10 inch color negative and then converted to digital data by using the 50 micrometer




Figure 12b — Landsat scene of Chesapeake Bay recorded on color film in the APL Image Processing Laboratory to identify familiar landmarks in the companion scene (Fig. 12a) from the Seasat SAR. The first three optical bands (0.45 to 0.69 micrometers) from the Thematic Mapper were used as blue, green and red exposures for approximately 10^6 picture elements. Although the Landsat database is of lower resolution than the Seasat database, many features are more easily distinguished since image quality is not degraded by speckle noise, elevation layover, or the finite antenna effects that hinder the interpretation of Seasat imagery. The digital data in Fig. 12b were provided by R. E. Sterner from tape archives in the APL Submarine Technology Division.

aperture of the P-1700 scanner. Three passes were made in succession with red, green, and blue filters to record color separations on digital magnetic tape. The Skylab image was subjected to a 50% reduction in scale, in the same manner as that described for the Seasat image, producing an image with a spatial resolution of approximately 55 meters. The LSI-11 microprocessor controlling the Comtal image processor was used to apply a two-dimensional linear interpolation to the Skylab image matching the spatial scale of the Seasat image.

A cross correlation of the two Nantucket scenes was accomplished with Fourier transform methods,

using software published by the IEEE Acoustics, Speech and Signal Processing Society⁶ and recently installed on the PDP-11/70 computer in the Image Processing Laboratory. Both the Skylab and Seasat images were further minified, reducing the spatial resolution to 100 meters. At this resolution, most of the Nantucket scene can be represented by 512×512 pixels, which is the upper limit on the size acceptable for processing with the IEEE software. Each image was Fourier transformed, the amplitude-normalized cross-product was formed in the spectral domain, and the complex result was subjected to Fourier inversion, producing a cross correlation (Fig. 16) of the




Figure 13 — Four frames were extracted from a 16 millimeter color filmstrip produced on the Dicomed D48C filmwriter depicting surface waves in the Atlantic Ocean as viewed by the SAR aboard the Seasat satellite in 1978. (a) The pseudocolor-encoded contours of the spectrum obtained for a region of moderate surface wind assist in the distinction of two surface wave systems generated by two separate storm centers. (b) The dominant system with a wavelength of approximately 200 meters can be discerned in the corresponding radar scene of the ocean surface. (c) The banding of wave power in the vertical dimension of the spectrum obtained for a region of high wind indicates a loss of resolution in the azimuth coordinate. The color tablet shows the relative spectral intensity of the color contours as linear steps of three standard deviations above the background noise, which is depicted in continuous shades of gray. The circles in the spectrum, from largest to smallest, represent wavelengths of 50, 67, 100, and 200 meters, respectively. The direction of wave propagation is referenced to the six angular grid lines in 30° increments relative to the satellite velocity vector, which appears in the margin along with a North arrow. (d) The loss of image resolution in the azimuth (or along-track) dimension is manifested as vertical streaks in the Seasat ocean scene corresponding to the region of high surface wind. The digital data for this study were purchased by the APL Space Department from MacDonald, Dettwiler and Associates, Vancouver, B. C.

Skylab and Seasat images. This procedure was also applied to the red and green separations of the Sky-lab image to obtain a reference cross correlation for two images that are identical geometrically.

Optical images of the earth are often used for the classification of terrain by defining pixel subsets within an image that are bounded by limits imposed on the scene reflectance. The intensity limits applied for pixel classification may be specified with regard to one or more images of the same region, with the classification becoming more precise as more spectral bands are employed. The first step in this process is the interrogation of the images to be used for classification to determine the most frequently occurring combinations of image intensity.

The Seasat L-band SAR image of Nantucket was used in combination with one of the Sky-lab color separations to see how many different classes of ter-

rain could be distinguished. When only two images are involved, a scatter plot can be developed by counting the number of pixels that occur for each combination of Seasat and Sky-lab intensities. The PDP-11/70 computer was used to generate a scatter plot for identifying Nantucket terrain classes (Fig. 17), using the Seasat image in combination with the Sky-lab green color separation. The Comtal image processor was used to display the scatter plot and add a pseudocolor code. Three terrain classes were defined as rectangular subsets within the domain of the scatter plot. Then the Comtal system was used to create a color display of Nantucket where the three major terrain classes would be represented by blue, green, and red. Minor classes (e.g., the town of Nantucket and its airfield to the south) have pixel counts so low that they do not show up on the scatter plot. However, they may be defined by interrogating the




Figure 14 — The Seasat SAR recorded this digital image of Nantucket Island in 1978. Note that strong specular reflectors such as the sandy beach and airstrip do not scatter microwave radiation back toward the side-looking radar. Multiplicative noise is apparent in a comparison of the wind-streaked seas to the south of the island with the calm waters in the sheltered cove. Surface manifestations of the shearing force created by the bathymetry in resistance to the local circulation are also evident in this digital scene processed by the Jet Propulsion Laboratory correlator and mapped to ground coordinates in the APL Image Processing Laboratory.



Figure 15 — This image of Nantucket was recorded by a camera aboard Skylab in 1976. The sandy beach around the island is a strong reflector of optical radiation in all spectral bands. The shelf along the north coast of the island and the shoals to the southeast are imaged by reflected sunlight that has penetrated the ocean surface. The original Skylab print was provided by the National Oceanographic and Atmospheric Administration and digitized in the APL Image Processing Laboratory.

Skylab and Seasat images in those regions to discover what intensity limits should be imposed. In this way, a fourth class was defined as high reflectance in both the L-band Seasat SAR and the green optical band of Skylab. This class is represented by white in the terrain classification map and includes large manmade structures such as the town of Nantucket at the western end of Nantucket Harbor. The airfield reflects optical radiation well, but its smooth surface does not backscatter L-band radiation toward the Seasat sensor. Hence, the airfield appears bright in the Skylab image and dark in the Seasat image. A fifth class could be defined as this particular combination of intensities, but it is not necessary since the

airfield is well defined in black, identifying those regions not included in any of the other four terrain classes.

Space Shuttle Imaging Radar

A synthetic aperture radar was carried aboard the space shuttle *Columbia* during its second flight. The remote imaging system was designed at the Jet Propulsion Laboratory (JPL) and was named the Shuttle Imaging Radar — Model A (SIR-A). Its primary goals during the NASA coordinated mission were measurement of the imaging system's characteristic response over land targets and determination of shuttle stability requirements during image acquisition. In addi-




Figure 16 — (a) The convolution product for the Skylab green separation versus the Skylab blue separation is shown as a test of the repositioning accuracy of the P-1700 scanner and as a reference for other geometric correlations. Note that the product is largest at the center and that statistical fluctuations correlate at a low level as one color separation is spatially shifted relative to the other. (b) Also shown is the convolution product for the Skylab green separation and the Seasat image. The central region indicates that the scenes are registered well and that both positive correlations (white) and negative correlations (black) exist for these two sensors. Random speckle noise degrades the Seasat intensity distribution so that accidental correlations occur at a relatively higher gray value.

tion, the imaging radar was switched on over selected ocean regions to record the surface effects of currents and of other large scale circulations.

The occurrence of internal waves in the Andaman Sea is a quite common phenomenon, so it is not surprising that the shuttle captured an image of one such wave system. A 4000×2000 pixel image, represent-

ing 100 kilometers in azimuth along the 50 kilometer swath, was output from the JPL digitizer and received at APL as a computer tape. The image intensity was observed to fall off rapidly at both extremes of the shorter dimension. That trend was approximated by an average over the lower half of the longer dimension and was removed from the entire image. After that trend removal had been accomplished, the data were read onto the disks of the PDP-11/70 computer with a Fortran program that reduced the image resolution by a factor of 2 in each dimension. The shuttle scene was then transferred to the Comtal Vision One/20 image processor, where minification by linear interpolation resulted in a 512×512 pixel data base. The PDP-11/70 computer was used to calculate an average intensity profile in the direction of wave propagation (Fig. 18), which was then subjected to one-dimensional Fourier analysis.

The space shuttle scene of the waves in the Andaman Sea was subjected to two-dimensional fast Fourier transform analysis using a software package installed on the 3033 computer.⁷ Removal of the image mean and windowing with a circular cosine function were accomplished on the PDP-11/70 computer in a preprocessing stage before a fast Fourier transform was applied, after which a two-dimensional power spectrum was computed. When the power spectrum was displayed on the video monitor of the Comtal image processor, it was obscured by a large-intensity spike near the center of the display. Therefore, the PDP-11/70 computer was employed to compute the natural logarithm of the power spectrum so as to enhance the low-intensity features nearer the periphery (i.e., toward larger wavenumbers). A false color code was applied depicting lower spectral intensities as blue, moderate spectral intensities as green, and the highest spectral intensities as red. The dominant wavelength for the internal wave train was identified, and the direction of propagation was determined.

In the next few years, another radar system (SIR-B) will be carried into space aboard the space shuttles for a survey of the oceanic currents near the polar regions of the southern hemisphere. Internal waves are common in that region as well, and it will be interesting to observe their interactions with wind waves and ocean swell.

Differential Radiography

Historically, in medicine, physicians have visually compared, analyzed, and interpreted planar projections of the anatomy. Many biological experiments and most radiographic examinations still rely on film as the recording, storage, and display medium for data distributions in two dimensions. Now, however, electron microscopes, digital fluoroscopes, and a variety of medical scanners are controlled by computers, and anatomic images are recorded on magnetic disks and tapes. Since these anatomic images are stored in a digital format, many clinical radiologists and biomedical research scientists have become




Figure 17 — This classification map of Nantucket was computed using the Seasat and Skylab scenes of the region. The insert in the lower left corner is a pseudocolor-encoded scatter plot depicting the cumulative pixel count for all combinations of Seasat and Skylab intensities. Clusterings within the scatter plot represent distinct classes of terrain that have been bounded by the intensity limits that are shown as white rectangles. The classification map was generated by assigning a distinct color to each class of picture elements. In this example, white, red, green, and blue pixels correspond roughly to geographic concentrations of buildings, sand, land, and water, respectively.

interested in digital image processing for contrast enhancement, feature extraction, and computer-assisted diagnosis. It is important that digital processing enhance the visual cues needed for diagnosis without sacrificing the general appearance of the medical image.

The computer facilities of the Image Processing Laboratory and its capabilities for biomedical image processing have been applied to enhance fluoroscopic arteriographs. Two fluoroscopic X-ray films of the abdomen were obtained from the Johns Hopkins Hospital to experiment with enhancement algorithms before a digital fluoroscope was acquired. The films had recorded the spleen and its arterial circulation

before and after the infusion of radio-opaque contrast material via a catheter advanced through the abdominal aorta to the celiac artery. The X-ray films were scanned individually on the P-1700 film digitizer after visual registration, to record two digital images on magnetic tape. These digital representations of the X-ray films were transferred to the Comtal image processor, where they were shifted horizontally and vertically to achieve optimum translational registration. The image recorded before the contrast material was injected (Fig. 19a) was subtracted from the image recorded afterward (Fig. 19b) in order to eliminate the bony features common to both. A histogram (pixel count versus pixel intensity) was gener-

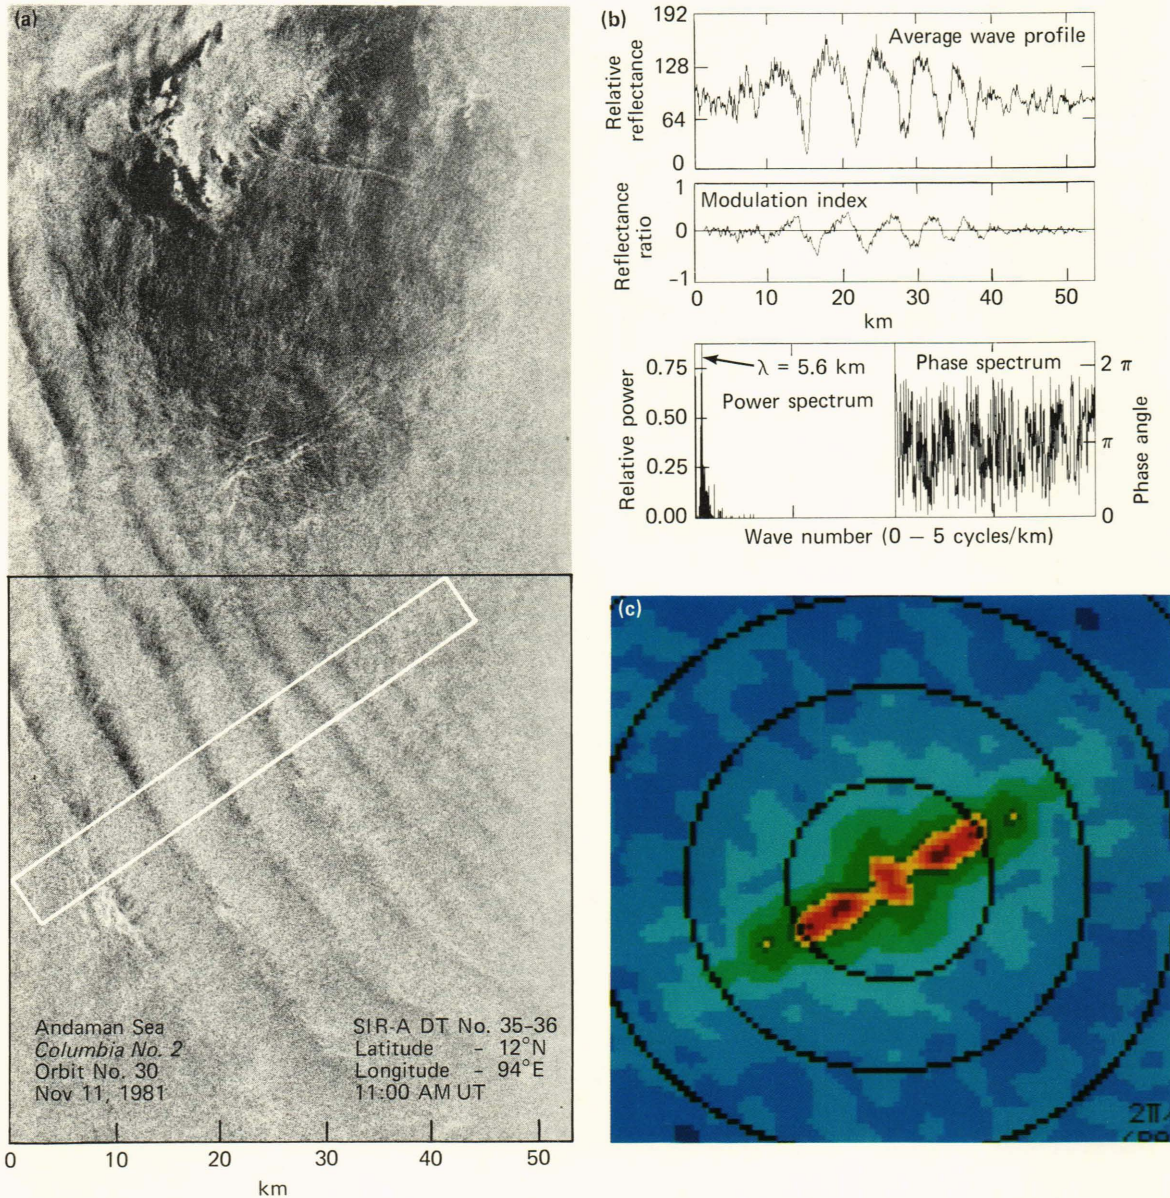


Figure 18 — Waves measuring approximately 5 kilometers between crests were recorded by an imaging radar carried in the cargo bay of the space shuttle *Columbia* during its second flight. Data enclosed by the rectangle (a) were used to compute the average wave profile and its modulation index (b), as well as the power spectrum and phase spectrum resulting from a one-dimensional fast Fourier transform of the average wave profile. A two-dimensional fast Fourier transform was applied to all of the image data to obtain a logarithmic power spectrum (c). A pseudocolor code has been applied to the spectrum depicting contours of the logarithmic wave intensity. These internal waves were detected in the Andaman Sea, where the regional circulation is conducive to their formation. The digital data base for this radar image was provided by C. Elachi of the Jet Propulsion Laboratory.

ated for the difference image (Fig. 19c), and a linear contrast stretch was applied to utilize fully the 8-bit range of the raster video display. It is expected that this type of differential radiography will become a standard procedure as software packages for image analysis are developed by vendors of digital fluoroscopes.

The difference image of the splenic artery was edge-enhanced using the image convolution and image arithmetic pipeline processors of the Comtal image processor. Convolution was applied to the differ-

ence image twice in succession to implement an effective 5×5 filter with the hardwired 3×3 pipeline processor. The result of the 5×5 convolution of the difference image was combined with the difference image to create an edge-enhanced image (Fig. 19d) using the Comtal arithmetic processor. The effect of depth perception (Fig. 19e) was simulated with edge shadowing, produced as a result of the spatial shift introduced during the convolution process. The two derivatives of the radiographic image were used as the green and blue components of a color display

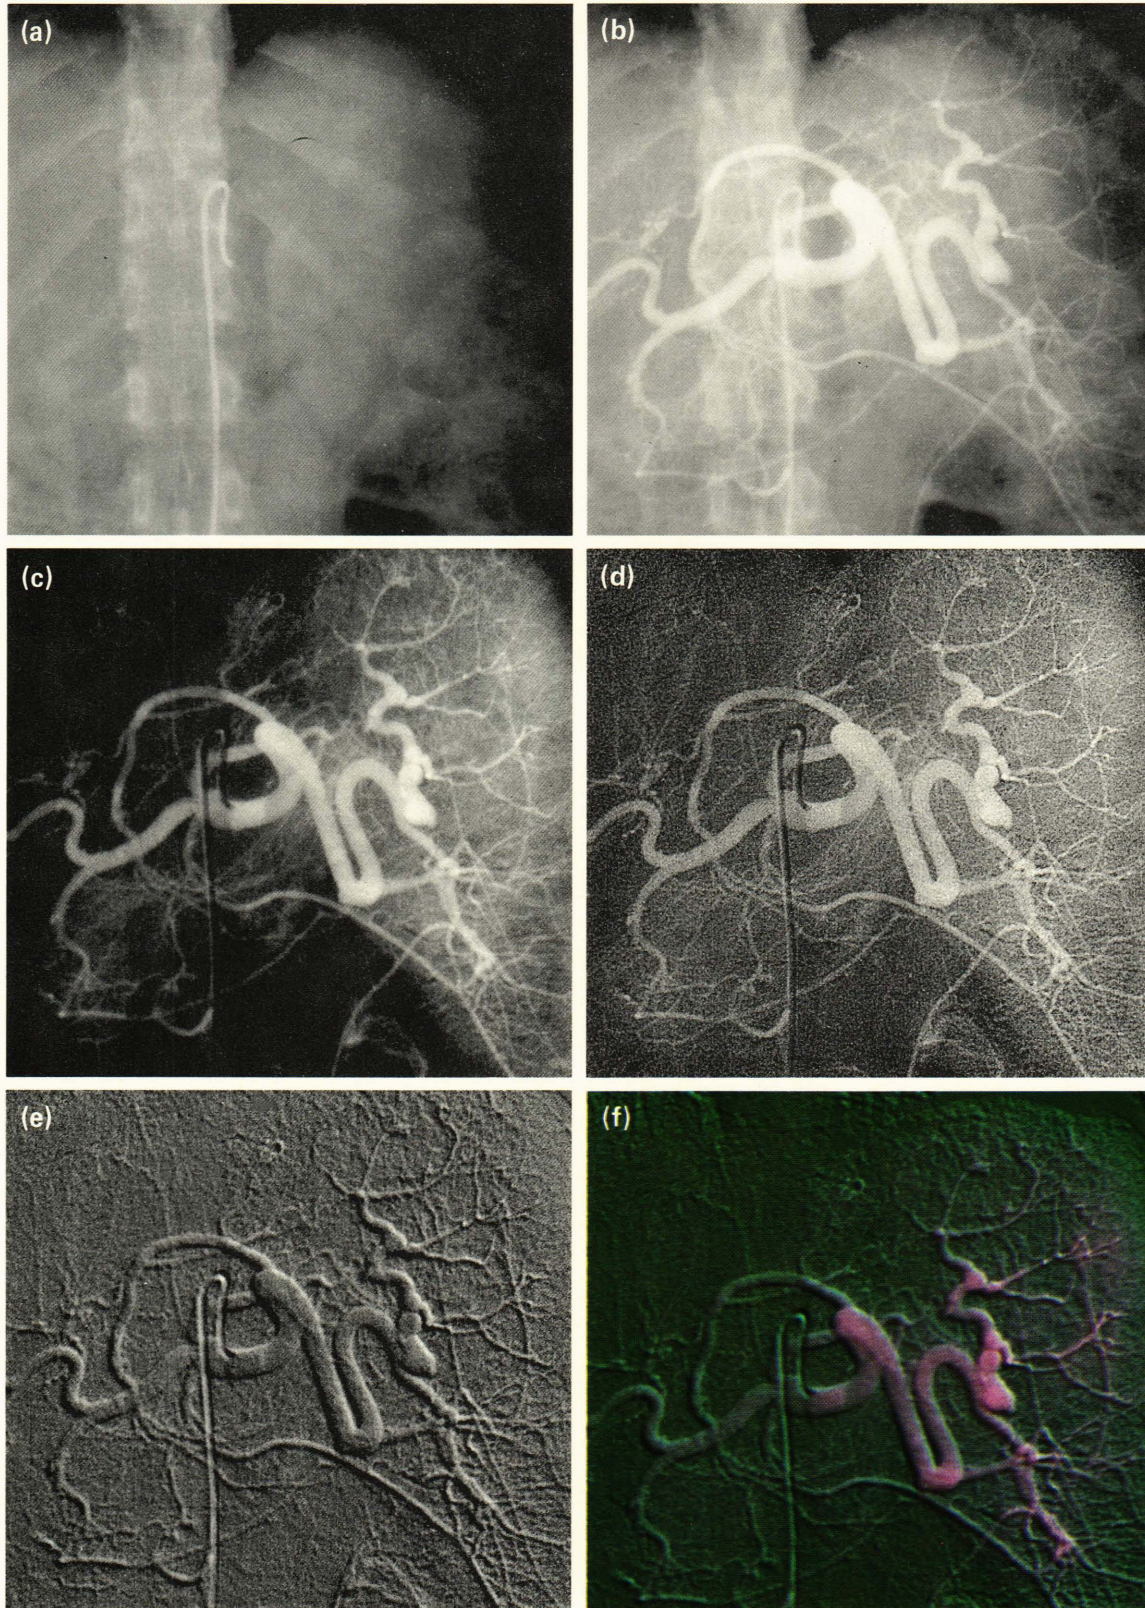


Figure 19 — The capillaries of the splenic artery supplying blood to the spleen have been emphasized in this fluoroscopic study showing the abdomen (a) before and (b) after infusion of radio-opaque contrast material. The difference image (c) was obtained by subtracting (a) from (b), with accurate registration to eliminate common features such as the spinal vertebrae and rib cage. Two methods of edge enhancement were applied to the difference image, and the results are shown (d) without shadowing and (e) with shadowing. The false color display (f) was created by using (c),(d), and (e) as red, blue, and green exposures on color film, respectively. The X-ray films for this study were provided by R. Chang of the Johns Hopkins Hospital.

with the original difference image used as the red component (Fig. 19f).

The entire process applied to the original pair of X-ray images, from the calculation of a difference image to the computation of an edge-enhanced image, can be executed with the Comtal system in less than 1 second. To achieve this speed, it is necessary to enter a sequential list of Comtal keyboard commands into the random access memory of the LSI-11 microprocessor. The list may be saved on magnetic tape or disk for future applications.

Edge enhancement is more difficult to standardize for the clinical assessment of radiographs since convolution and image combination may be defined in a number of different ways. Furthermore, the optimal choice of an edge-enhancement algorithm may depend on the size and shape of anatomical features or the spatial resolution of the medical image. In spite of these complications, the edge enhancement of biomedical data can be helpful in resolving small and intricate features. Therefore, manufacturers of computer-controlled systems for the acquisition of biomedical images are likely to provide general edge-enhancement algorithms as analytic aids to clinical diagnosis.

Computerized Tomographic Scan of the Brain

The use of image processing technology by medical researchers is made more productive by software systems that promote the direct interaction of the researchers with their images. The development of such a user-friendly environment for certain types of analysis has been the object of a project involving the Grinnell image processing system.

The Grinnell system has been used to develop programs to distinguish between tissue and fluid in the brain (Fig. 20). The programs will be used to compare brain scans of normal individuals with those suffering from major psychoses such as schizophrenia and manic-depressive illness.

The functions provided by the programs carry out image enhancement and allow the clinician to outline interactively regions of potential interest. These regions can then be investigated further through texture analysis such as two-dimensional autocorrelations or through generating simple point statistics such as means and standard deviations.

The routines are written in the APL-language implementation of the Grinnell Command Language. The APL environment allows quick development of algorithms and programs without the distractions of compilation or linking. This is especially important for users whose main interest is in the algorithms and techniques rather than in the programming. The Grinnell Command Language exploits these features of APL to allow the algorithm developer direct control of the hardware from an interactive high-level language.

Guideway Movie

Automated Guideway Transit Systems represent a relatively recent development in transportation tech-




Figure 20 — A set of programs developed by A. D. Goldfinger and G. A. Heyler for the Grinnell system attached to the 3033 computer to allow clinicians to select and analyze regions of the brain interactively. The techniques developed will be used to compare brain scans of normal individuals with those of individuals suffering from major psychoses such as schizophrenia and manic-depressive illness. The pseudocolor-enhanced brain scan shown here was generated by computer-assisted tomography. This work was done for G. D. Pearlson of the Johns Hopkins Medical School.

nology and are characterized by the operation of small- to moderate-sized vehicles under automatic (driverless) control on dedicated guideways. Such systems will be required typically to support more vehicles than do conventional transit systems in order to attain satisfactory levels of capacity, thus necessitating close regulation of vehicle speeds and spacings to assure safe, comfortable, and reliable operation.

Research on Automated Guideway Transit Systems that employ vehicle-follower control has been conducted at APL.⁸ Under vehicle-follower control, the motion of an individual vehicle on the guideway is controlled by the motion of the preceding vehicles. Part of the work done on guideway controllers was the development of a discrete network simulation. The main elements of the network were a merge link on which vehicles entered the guideway, a station that vehicles entered to load or unload passengers, the guideway itself (a rectangular track), and an exit link on which the vehicles left the guideway. Of particular interest were the two regions where the merging of vehicles would take place and the station where vehicles would stop for passengers.

The network simulation dealt with such dynamic events as the merging of vehicles onto the guideway.

Such events are readily displayed by animation. Accordingly, the discrete vehicle simulation was used to generate an animated film illustrating the operations of the guideway controller.

This computer animation is a 9-minute color film that illustrates all the important elements of the digital simulation, such as the vehicles merging onto the guideway, stopping to load and unload passengers, and exiting the guideway. In addition, the movie features a number of special animation effects such as a dissolve from the title sequence to the guideway sequence and the display of closeup views of the two regions of the guideway where the merging of vehicles takes place. Figure 21 shows two frames from the movie.

The Guideway film was produced on the Laboratory's 3033 computer using the DISSPLA graphics software and output on the Dicomed film recorder. Since DISSPLA allows graphics to be created with high-level commands, special features such as fades, zooms, and high quality annotation can be included in animation without great programming effort.

The computer program, which used DISSPLA to create the graphics, was a modular PL/I program driven by the output file from the discrete vehicle simulation program. The output file contained the positions of all vehicles for each frame in the movie, together with other information such as the velocities of the vehicles. Since changes were made in the movie throughout its development, particular attention was given to making the program flexible. It was structured with individual subroutines, corresponding to the graphical elements that were to appear in the film. For example, separate routines were written to draw the vehicles, the guideway, the clock, the shaded background, and so on.

To exploit this flexibility, a script file was set up as part of the input to the program. The script file was used to select the sequence of frames produced, the graphical elements to appear in the frames, and the device on which the graphics were to be generated. This was particularly effective for debugging the program and for tuning parameters that controlled motion, colors, and the positions of graphical elements in the frames. By changing the script file, it was possible to examine any set of frames interactively on the IBM 3279 color graphics terminal, to generate color slides of individual frames, to generate movie sequences containing only selected graphical elements (e.g., for example, only the vehicles and the guideway to check motion), or to generate a complete sequence for the final version of the movie on the film recorder.

An important graphics feature supported both by the DISSPLA software and by the software that controls the Dicomed film recorder is picture segmentation. With segmentation, static graphical elements can be defined once by the program running on the 3033 computer and then reused as often as necessary when the film is generated on the film recorder. Most of the vectors that appear in any frame of the guide-




Figure 21 — A 9-minute color computer-animated film was made to show the movement of vehicles on an Automated Guideway Transit System that uses vehicle-follower control. These two frames from the movie illustrate two different cases of vehicles merging onto the guideway. (The velocities of the vehicles are indicated by their colors, with green representing maximum velocity.) The top frame shows the action of the guideway controller in effecting smooth merges of vehicles entering the guideway. The bottom frame illustrates the merging of vehicles leaving the station.

way movie are unchanged for a large number of frames. Therefore, the use of segmentation greatly decreases the running time of the Guideway program on the 3033 computer. (Segmentation causes a slight increase of the running time on the Dicomed since segment definitions must be retrieved from disk each time the segments are referenced.) The Guideway movie shows that computer animation can be an effective technique to illustrate scientific research.

FUTURE PLANS

A strong, modern, central Image Processing Laboratory must be available to serve the needs of the numerous APL programs having requirements in the rapidly expanding field of image processing. To that end, a two-phase plan has been prepared for upgrad-

ing the Image Processing Laboratory over a period of several years.

The first phase of the plan has been approved and is currently being implemented. Two high-speed tape drives and two 300-megabyte disk units have been added, along with a line printer and the eight-pen plotter previously mentioned. Other equipment to be added includes an interactive computer-aided presentation graphics system, a back-lighted digitizing table, an electrostatic hard copy device, and a video digitizing system. The last item will permit digitizing for computer entry imagery transmitted as a video signal or recorded on video tape or disk.

The second phase of the upgrade will begin in 1984. Some of the items being considered are discussed below.

Eventually it will be necessary to replace the PDP-11/70 computer with another computer or to augment it with a second computer. The new computer would be a 32-bit machine with an addressable memory large enough to store a full image, allowing fast "in core" algorithm techniques. The operating system would support virtual memory, freeing programmers from the size constraints of the current program. An array processor will be considered as a means to increase computing power, with or without another computer. The vector processing capability of such a device could improve the throughput rate 100 to 200 times.

Additional image processing work stations will almost certainly be required. Some of them may be more sophisticated than the present ones.

One or more high-resolution displays will be considered to accommodate images of 1024×1024 pixels or larger. The present 512×512 pixel displays are adequate for most work, but some tasks, such as the display of Landsat imagery, require more resolution.

Charge-coupled device image sensors hold considerable promise for future image processing applications, and optical disks offer the possibility of storing up to 4×10^{10} bits of imagery on a single disk with a data rate of 3 megabytes per second.

REFERENCES

- ¹H. W. Ko, J. W. Sari, and J. P. Skura, "Anomalous Microwave Propagation Through Atmospheric Ducts," *Johns Hopkins APL Tech. Dig.* **4**, 12-26 (1983).
- ²D. M. Silver, "Rotationally Inelastic Collisions of LiH with He.I. Ab Initio Potential Energy Surface," *J. Chem. Phys.* **72**, 6445-6451 (1980).
- ³R. C. Beal, D. G. Tilley, and F. M. Monaldo, "Large and Small Scale Spatial Evolution of Digitally Processed Ocean Wave Spectra from SEASAT Synthetic Aperture Radar," *J. Geophys. Res.* **88**, 1761-1778 (1983).
- ⁴L. Fu and B. Holt, "Seasat Views Oceans and Sea Ice with Synthetic Aperture Radar," JPL Publ. 81-120, Pasadena, Calif. (Feb 1982).
- ⁵J. C. Curlander and S. N. Pang, "Geometric Registration and Rectification of Spaceborne SAR Imagery," International Geoscience and Remote Sensing Symp., (Jun 1982).
- ⁶D. Fraser, "Optimized Mass Storage FFT Program," in *Programs for Digital Signal Processing*, IEEE Press, New York (1979).
- ⁷J. G. Moik, *Small Interactive Image Processing System (SMIPS) User's Manual*, NASA-GSFC X-650-73-273 (May 1976).
- ⁸A. J. Pue, H. Y. Chiu, and S. J. Brown, Jr., *Operational Concepts and Implementation Techniques for Vehicle-Follower Control of AGT Systems*, JHU/APL CP 074/TPR 041 (Aug 1979).

ACKNOWLEDGMENTS—Laboratory projects described in this article have been funded by various contracts administered by the Department of the Navy, the Department of Defense, the Department of Transportation, the National Oceanographic and Atmospheric Administration, and the National Aeronautics and Space Administration.

The authors wish to thank the many people who have contributed to the image processing applications discussed herein by providing the tasks and the support, and in many cases doing much of the work themselves. In particular, we acknowledge the contributions of A. D. Goldfinger, J. P. Skura, J. R. Apel, R. E. Stovall, S. A. Hartson, D. K. White, H. W. Ko, R. C. Beal, G. D. Smith, and K. S. Bonwit.

Finally, we thank H. J. Reider, W. E. Hanson, J. R. Sanchez, and J. G. Schmidt for the production of the imagery used in the illustrations and A. J. Mancuso and S. A. Mitchell of the Central Photographic Laboratory for the color film processing. The assistance of J. A. Hagen and J. W. Stevens in maintaining the Image Processing Laboratory equipment is also greatly appreciated.