

THE IMAGE PROCESSING MARKETPLACE

R. C. Gonzalez
Perceptics Corporation & University of Tennessee
Knoxville, Tennessee

Invited Paper

HISTORICAL PERSPECTIVE

Interest in gray-scale digital image processing techniques dates back to the early 1920's, when digitized pictures of world news events were first transmitted between New York and London via the Bartlane cable picture transmission system. Pictures were coded for cable transmission and then reconstructed at the receiving end by specialized printing equipment. Use of the Bartlane system reduced the time required to transport a picture across the Atlantic from more than week to less than three hours.

Although improvements on processing methods for transmitted digital pictures continued to be made over the next thirty-five years, it took the combination of large-scale digital computers (the technology) and the space program (the need) to bring into focus the elements of what we refer to as modern digital image processing.

Work on using computer techniques for improving images from a space probe began at the Jet Propulsion Laboratory in 1964, when pictures of the Moon transmitted by Ranger 7 were processed by a mainframe computer to correct various types of image distortion inherent in the on-board television camera. These techniques served as the basis for improved methods used in the enhancement and restoration of images from such familiar programs as the Surveyor missions to the Moon, and the Mariner series of flyby missions to Mars. The digital image processing systems in these early applications were used primarily for image display, and were based on rotating-disk technology which was, at the time, the most economical and practical way to achieve the data transfer

rates needed to refresh a TV screen at 30 frames/sec.

The next major step in image processing hardware took place in the early 1970's, with the appearance of new image processing system designs based totally on solid-state technology. One of the keys to this technology was the availability of 16K RAM's that were both inexpensive and fast enough to allow in-memory storage and display of large (512 x 512 x 8 bits) images at 30 frames /sec. This, along with the availability of relatively inexpensive minicomputers such as the DEC PDP-11 series, brought complete image processing capability to a large number of users in academic and research institutions. This new breed of users made significant contributions to the field, and provided the theoretical and applied foundation of image processing through the 1970's.

Although the memory address space of early minicomputers and peripheral storage media for these systems were grossly inadequate by today's standards, a number of new application areas with commercial potential began to appear, ranging from medical imaging to remote sensing. This growth in the field was further fueled by the appearance of "superminicomputers," such as DEC's family of VAX computers. These machines, with their virtual addressing architectures, were ideally-suited for data intensive applications involving large data arrays such as those found routinely in digital image processing. When coupled with supporting peripheral equipment (principally large magnetic disks) and image processing systems with capabilities for on-board hardware processing, virtual machines provided a processing environment which was the

standard throughout the mid 1970's and early 1980's. In spite of continually declining prices, however, systems like these remained expensive enough so that their use was focused primarily in government-related applications (with the Department of Defense being the lead user) as well as in academic institutions and research laboratories.

In the early 1980's, new image processing products began to appear in the form of individual boards that plugged into a standard bus and thus allowed a user to integrate image processing capability directly into a host computer. During this period, the cost of suitable host computers dropped dramatically with IBM's introduction of the personal computer in 1981. IBM's stamp of approval on these low-cost processors, and the fact that the PC had an open architecture, spawned dozens of start-up operations dealing with the design and manufacturing of image processing board sets for the PC. Improvements of these boards, and subsequent introduction of the IBM PC AT, the Sun engineering workstation and, more recently, the Macintosh II, have had a dramatic impact on the cost of complete image processing systems. Today, one can assemble a desktop system for \$20,000, with the same capabilities (and a fraction of the size) of systems that less than a decade ago cost more than \$100,000.

As a result of these developments, the "micro-imaging" market appears to be on the verge of explosive growth, with electronic publishing, medical imaging, biological sciences, and industrial inspection leading the way. Even by conservative estimates, the average, compounded annual growth in these fields is expected to exceed 50% a year.

THE CURRENT MARKETPLACE

The largest current applications of commercial image processing are machine vision, printing & publishing, graphic arts, medical imaging, and remote sensing. This is in addition to the well-established academic and government R&D markets for image processing products and services.

Machine vision deals with applications typically characterized by autonomous processing. The largest exponent of this technology is automated industrial inspection, but security and surveillance are also significant and growing rapidly.

Electronic printing and publishing uses digitizing image scanners to capture photographic images, text, and line drawings. Image processors enhance, size, orient, and modify these images, and merge them with text files for inclusion in documents and other publishing formats. Typical uses of image processing in electronic publishing include pasteup, cropping, editing, page composition, screening, and related special effects.

Graphic arts principally involve package design and poster art. Graphic design arts traditionally combine image processing techniques with computer-generated images to produce attention-getting visual products. Typical applications include package design, advertising art, and artistic posters.

Medical uses of image processing consist of diagnostic, clinical, and research applications. Diagnostic medicine is a large user of image processing products as the basis for CT, MRI, digital radiography, and nuclear medicine.

Clinical uses of image processing are also growing rapidly. One area of potential large-scale usage of image processing is in microscopy involving automated or semi-automated measurements and counts of microscopy samples.

Remote sensing deals with image acquisition of the Earth's surface using satellite- and air-borne imaging sensors. This market segment provided one of the first commercial applications of image processing technology. This is still a major market, with applications that include terrain mapping, resource exploration, and marine fish location, in addition to numerous military uses.

A recent study by Frost & Sullivan (New York) estimates the 1988 U.S. market in these fields to be on the order of \$600M, a figure that is expected to more than double in the next three years. Dataquest Inc. (San Jose, CA) predicts an annual compounded growth rate for desktop image processing units in the above fields to exceed 50% through the end of this decade.

STATE OF THE TECHNOLOGY

The principal functions in image processing are sensing, preprocessing, segmentation, recognition, interpretation, display, compression, transmission, and storage. The principal sensing modalities for digital image acquisition are area arrays (typically in the form of TV cameras) and hardcopy scanners. The state of the art in monochrome TV cameras is pushing 1024 x 1024 spatial resolution with 8- to 10-bits of intensity resolution; color cameras with 512 x 512 (8 bits per plane) are now becoming available. Hardcopy scanners with 400 dots-per-inch and higher are common-place.

Preprocessing deals with functions such as noise reduction, enhancement, and restoration. This is one of the most technologically advanced segments of the image processing field, primarily because preprocessing functions generally lend themselves well to theoretical formulations.

Segmentation deals with the (automatic) partitioning of an image into its constituent parts or objects. This is one of the most important steps in the eventual formulation of automated scene analysis procedures. In most commercial applications of image processing, segmentation is usually based on single-threshold methods or simple adaptive techniques. However, this is far from being an unsolved problem, as evidenced by the hundreds of articles on this topic that have appeared in the image processing literature in the past five years.

Recognition is basically a labeling process that assigns an identity to each object or other element of interest resulting from segmentation. Interpretation is the process that assigns meaning to an ensemble of labeled objects. Both of these topics deal with automated image analysis and, when compared with the capability of humans for performing similar tasks of perception, are basically in a rather primitive state of development. Commercial implementation of recognition and interpretation techniques are generally limited to matching techniques (either by correlation or structural formulations).

The state-of-the-art in image display is on the

order of 1024 x 1024, with 8 bits of intensity resolution and a raster refresh of 60 Hz, noninterlaced. Both monochrome and color displays with these capabilities are readily available commercially.

Image compression is used to reduce transmission bandwidth and storage requirements. Due to the relatively large data arrays typical in this field, image processing has traditionally been a premier user of the latest techniques in data compression and storage. In terms of data compression, it is not unusual today to achieve an information-preserving compression ratio of 4-to-1, and a non-preserving ratio of 32-to-1 for moderately complex gray-scale images.

The "dream" of transmitting images quickly and cheaply via low grade telephone lines dates back to early stages of development of the picture phone. Although advances in data compression have helped the transmission bottleneck for digital images, this continues to be a topic of considerable commercial interest. Recent advances in optical transmission lines, and the ever-increasing availability of satellite links, have made it economically feasible to transmit image data at rates that are becoming practical. For instance, it is now feasible to transmit medical and news imagery anywhere in the world at prices that are justifiable when compared to the business gains achieved by the fast transmission of pictorial data over commercial communication networks.

Recent advances in optical (laser) disk digital storage media are beginning to have an impact in the image processing field. "Write-Once-Read-Many" (WORM) laser disks with the capability to store 2 GBytes of digital data in a single platter are readily available, and "jukeboxes" capable of storing and retrieving over 200 Gbytes of data have appeared in the marketplace during the past 18 months. Added to the fact that this technology has a shelf life in excess of 30 years, laser disks are beginning to replace magnetic tape as the medium of choice for archiving digital image data.

In terms of hardware technologies for executing image processing algorithms in the areas just discussed, the state of the art offers a variety of choices ranging from board-level products and desktop workstations to large

scale image processing systems interfaced to mainframe computers. It is of interest to note that most of these products achieve acceleration via the use of dedicated hardware (e.g., pixel processors) based on pipeline architectures. Parallel processing, a topic discussed in more detail in the following section, has just begun to have an impact on image processing products.

A GLIMPSE AT THE FUTURE

Although the technologies discussed in the previous section will continue to fill new market opportunities in the foreseeable future, the single area likely to have the most fundamental impact on the breadth of future markets is automated scene analysis. This will require a leap in image segmentation, recognition, and interpretation technologies.

Improvements in segmentation are likely to come from work dealing with new sensing methods developed either to exploit the inherently three-dimensional nature of scenes, or to extract directly other attributes such as color properties. For example, relatively new approaches that combine range and intensity data dramatically simplify object boundary detection, a classic problem in 2D image processing. This approach is an example of a broader class of problems dealing with multisensor fusion, whose principal objective is to simplify image segmentation.

Once an image has been segmented, we are still faced with the problem of describing its segments in a form suitable for computer processing. In general, a "good" description often has the potential for making recognition of individual image segments a trivial task. However, even in 2D processing, descriptors that are independent of relatively simple geometric distortions are difficult to formulate. When this is carried out into 3D representations, the problem becomes even more severe. One approach likely to offer a partial solution to this problem is the use of knowledge-based modeling techniques designed to incorporate "expected scenarios" into the image analysis process. Only through effective techniques for limiting the operational environment of an image analysis system are we likely to develop description methodologies that are rugged enough for practical solutions.

Improvements in recognition and interpretation techniques are likely to emerge from the development of methods for incorporating knowledge into image analysis systems. One of the principal factors affecting the commercialization of these systems is the lack of generality achieved in the solution of any one problem. It is not unusual to spend several man-years in developing a system for a specific application, and then finding that the resulting technology is only fractionally transportable to the solution of a different problem. Given the state of the technology and the poor likelihood that any general solution to this problem will be found in the foreseeable future, a meaningful course of action is to focus attention on procedures that simplify system design. For instance, great promise is being shown by approaches that start with situation knowledge in the form of expert rules and use this information to produce decision trees ready for execution in a target machine. This type of approach has the potential for reducing the design time needed to produce a complex pattern classifier from years to months, and has the added advantage that maintenance issues are simplified to very attractive levels.

It should be evident that the issues just discussed have extensive computational implications. As mentioned at the beginning of this section, parallel computer architectures offer a potential solution to this problem. A particularly attractive implementation is the hypercube architecture because of its generality. Significant progress has been made in the past few years in the development of massively parallel hypercube systems. In fact, systems are now available (at a fraction of the cost of traditional supercomputers) that have the computational power of a thousand VAX 11/780 processors working in parallel. These systems occupy less than a cubic meter in space and can be cooled by traditional forced-air techniques. Although the cost is still high (on the order of \$2M) the fundamental approach of using thousands of inexpensive processors (as opposed to a few super-fast processors that have to be cooled cryogenically) is important for two reasons. First, it is likely to lead to cost-effective supercomputer capabilities in the not too-distant future; and second, the development of parallel computational algorithms will introduce a "think parallel" approach in the formulation of higher-level scene analysis solutions.