

Machine Vision in the 1990s: Applications and How to Get There

D. Petkovic

IBM Almaden Research Center, San Jose, California

J. Wilder

Rutgers University, New Brunswick, New Jersey

Introduction

This report contains a collection of papers from participants in the panel discussion “Machine Vision in the 90s: Applications and How to Get There,” which was held during the International Conference on Pattern Recognition, in Atlantic City, N.J., June 20, 1990. The participants were Dr. M. Ejiri (Hitachi), Prof. R. Haralick (University of Washington, Seattle), Prof. R. Jain (University of Michigan, Ann Arbor), Dr. P. Ruetz (LSI Logic), Prof. J. Sklansky (UC Irvine), and C.W. Swonger (Director, Image Processing, ERIM, Ann Arbor). Panel chairmen were Dr. D. Petkovic, IBM Almaden Research Center, San Jose, CA, and Prof. J. Wilder, Rutgers University, New Brunswick, NJ.

My motivation in organizing this panel was as follows. After many years working in the fields of machine vision and pattern recognition, I had certain observations that concerned me:

- Our basic understanding of the process of visual information analysis (recognition, measurement, etc), be it by humans or machines, is still poor. I think this is the main reason that it is so hard to construct machines that can analyze visual information effectively (which is the task of machine vision). It is extremely difficult to adapt these machines for a variety of tasks and almost always we have to start from “scratch.” The situation is not much better even in the restricted domains of industrial machine vision. This is quite contrary to graphics (a process of image synthesis), which is well understood and successfully implemented in practice.
- It is very difficult to model and predict the performance of machine vision systems. The potential rewards of purely theoretical work are limited,

and we need to perform costly large-scale experiments.

- The advances in these fields for the last 10 to 20 years have mainly been in technology (sensors, processors, memory, storage) and not in algorithms. Technology is still advancing with the same speed, whereas our basic understanding of vision is not.
- The transfer of research ideas into practice has been rather slow.
- Companies whose main product is machine vision equipment are not doing well. Industry is scaling back on full automation. The job picture for machine vision practitioners is not very bright.
- At the same time we see tremendous growth in computing, which has spurred growth in the use of image (office applications, business, insurance, etc.).
- Machine vision has to have a strong commercial base in order to survive as a scientific/engineering field in the long run.

When one looks at all this, the natural question is how to fulfill the potential of machine vision and what are the factors preventing this from happening. Therefore, we presented the following questions to the panelists:

- What are the new applications for machine vision in 1990s?
- What are the old applications that are now possible due to advances in technology?
- What research is needed to help make this happen?
- What are the technologies needed to help make this happen?
- How can the transfer of research ideas into practice be accelerated?

The emphasis on applications comes from my belief that, while it might take many years to understand human vision and general vision, the “next best thing” we can do is to devise machines that can do certain tasks (i.e., applications) cost effectively. These machines may not be too general, nor will they necessarily resemble human vision. If we are able to solve important problems for industry, society, and government, the field of machine vision will progress.

The panel was structured as follows. Each panelist presented his views (motivated by the preceding questions, but not limited to them) for about 10 minutes, and then we had a general discussion with lively audience participation. The panel lasted 2 hours and approximately 200 people attended.

This report contains a collection of short papers written by panelists. They are essentially “position papers” and represent, in my opinion, excellent reading for researchers, practitioners, strategists, and managers in this field. Rather than summarizing the main points, I recommend reading all the papers. I would just comment that if we have in mind benefits of the relentless progress in technology, and if we carefully guide our research, paying attention to the needs of the users, the field of machine vision will have a bright future.

I would again like to thank Dr. J. Sanz and Prof. H. Freeman for the opportunity to organize this panel. I would like to express my gratitude to the panelists who found the time in their busy schedules to participate and write a short paper. I would also like to thank my secretary W. Clayton for helping in the organization of this panel.

Dr. D. Petkovic
 Manager, Advanced Algorithms, Architectures
 and Applications
 IBM Almaden Research Center, San Jose, CA.

As co-chairman of the panel, I added my comments on the future of machine vision to those of Dr. Petkovic and raised some additional questions for the panel. My remarks were directed toward the focus of my work for a number of years—the industrial applications of machine vision. These comments and questions are summarized below.

Within the realm of potential applications of machine vision in industry, the largest percentage of human resources are currently devoted to on-line inspection and measurement tasks. It has been possible to automate some simple examples of these tasks using machine vision. Examples are found in the verification of the presence, position, and orien-

tation of objects, and the completeness of assemblies. Simple gauging operations have also been automated. However, many flaw detection and measurement tasks are currently beyond the capabilities of machine vision. Machine vision systems have not been robust enough, fast enough, or sufficiently cost effective to meet the requirements of high-speed industrial processes. But it is in just such areas that the need for machine vision is greatest since quality and safety frequently demand the highest standards of accuracy and consistency. Examples of such processes include the inspection of surgical needles, razor blades, and contact lenses.

Many of the limitations of current systems can be found at the front end, that is, in data acquisition and low-level image processing. High-contrast, high-resolution, noise-free images present few problems for machine vision systems. On the other hand, low-contrast, blurry, noisy images are frequently impossible to analyze. Some of the problems of data acquisition and low-level image processing are outlined here:

- A custom-designed optics and illumination system is required for almost every application. Active lighting is frequently required.
- A serious data-acquisition bottleneck exists in many high-speed applications. The speed of industrial processes is constantly increasing, speeds of 500 to 1500 objects/min in discrete processes and 1000 ft/min in continuous (web) processes are now common. In such cases massively parallel image processing systems are of little use if the visual sensors cannot acquire data fast enough to supply the system with the required information. (One approach to reducing the data bottleneck is to emulate the human eye by scanning at high resolution only in those portions of the image that are of interest and at decreased resolution elsewhere. Such a sensor is currently under development at the Rutgers University Center for Computer Aids for Industrial Productivity.)
- A high-speed, low-cost, compact, rugged three-dimensional camera is not available for those applications where triangulation and stereo techniques are not applicable.
- Robust algorithms are not available for extracting subtle flaws from highly textured surfaces and from complex shapes.
- Hardware for low-level image processing is still too expensive, for example, for pipelining multiple morphological operators. (Recent announcements from the manufacturers of vision hardware suggest that cost-effective solutions are becoming available).

The limitations on current data acquisition and low-level image processing techniques discussed earlier suggested the following additional questions that were posed to the panel:

- Should effort be expended on research leading to the development of a general purpose adaptive front end for machine vision systems? Conversely, are we constrained for the foreseeable future to continuing with customized designs of systems of optics, illumination, and sensing?
- Is it worthwhile to attempt to design inspection systems based on the scanning, local feature extraction, and scene analysis techniques of the expert human inspector? In other words, can the study of human visual processes lead to the development of better vision systems?

- Is the vision research community focusing sufficient effort in areas that will lead to robust, fast, cost-effective solutions to the challenging problems of industrial visual inspection and measurement?

A number of thoughtful responses by the panelists to these questions can be found in the papers that follow.

I would like to join my co-chairman in thanking Dr. Sanz and Prof. Freeman for their support and the members of the panel for their contributions to a stimulating discussion.

Prof. J. Wilder
Rutgers University, New Brunswick, NJ

Machine Vision in the 1990s

Masakazu Ejiri

Central Research Laboratory, Hitachi, Ltd., Kokubunji, Tokyo 185, Japan

1 Introduction

Since its inception "machine vision" technology has a history of almost 20 years. In that time a fairly large number of industrial processes have been automated by applying machine vision, most prominently in the semiconductor industry where demands have been enormous because of the precision and efficiency needed for production (Ejiri 1989b). However, there are still many processes that remain unchanged due to the difficulties encountered in securing the recognition reliability of machine vision.

It seems that "computer vision" research is rather methodology oriented and does not suggest pragmatic solutions to such difficult application problems that present themselves. That is, research and development seem to be saturated and stagnant from the application point of view. To open a new era in machine vision research, extensive discussions will thus be needed. This paper briefly discusses a few items that will be important for future machine vision applications.

2 The State of the Art of Industrial Machine Vision

There are two types of technologies in machine vision: passive and active. Machine vision systems that have been widely used in industry to date are based on a few simplified, fundamental methods. These are passive, monocular, two-dimensional-type vision methods including:

- Pattern matching method
- Feature parameter method
- Window method

The active technology typically utilizes a structured light to convey the information of the distance between the object and imaging device. An example of the structured light method is:

- Slit-light method

which is based on a triangulation principle.

These four methods are effectively used in industrial applications, sometimes in a singular form and

sometimes in a combined form (Ejiri 1989a, Sakou et al. 1989). The passive technology in three-dimensional-type vision, which typically includes the so-called shape-from-X methods where X implies such words as shading, texture, motion, and so on, is not yet fully applied to practical industrial use due mainly to the restriction from processing speed and reliability.

The technology level so far attained in machine vision applications can be summarized as follows:

Recognition of geometric features, supplemental features, and a few restricted cases of qualitative features of rather simply shaped objects. These objects must be within a predicted area or be positioned in a predetermined area. Usually the illumination is carefully controlled.

At present, machine vision techniques cannot be applied to:

Recognition of most of the qualitative features of even simply shaped objects and the geometric and supplemental features of complex objects, especially when the illumination drastically changes such as in outdoor, day-and-night use in all-weather conditions.

3 Possible Future Directions of Industrial Machine Vision Research

A variety of products based on different specifications will be needed in the future to meet the increasingly diversified needs of individuals. Due to this trend, the following two techniques are becoming important in future manufacturing research:

- Product realization technique
- Intelligent manufacturing control technique

The former is a design technique starting from the conceptualization of a product based on human needs and finishing with a product design through quick prototyping and evaluation, where the intelligent use of computers will be the key to success. The latter is a manufacturing technique of the designed products, where machine vision will play an important role.

In the manufacturing process two types of vision systems will be required: one to check continuously each manufacturing machine to see that the machine is kept in operation at its maximum efficiency; the other to check every product to see that it is being correctly manufactured. This inspection function must be furnished in a distributed manner at every important stage in the manufacturing process. In the 1990s the main stream of machine vision applications will thus be changed from assembly to inspection. In particular, inspection techniques that adapt to a group of various products produced under different specifications will be of utmost importance.

To achieve intelligent control of the manufacturing process, the following three items will be increasingly important in relation to the future machine vision research.

3.1 Sensor and Visualization/Enhancement Techniques

As for the sensors, sensitivity, resolution, and dynamic operation range are the major concerns for their applications. Machine vision practitioners have been forced to struggle with illumination problems in every application. Thus, the present MOS- and CCD-type image sensors, and their future successors, will not satisfy them. A new image sensor with an extremely wide dynamic range must be developed. The lack of dynamic range in sensors currently being used is making image processing extremely difficult. If a wide dynamic range image sensor were available, a fairly large portion of the image processing (mainly preprocessing) algorithms would not be necessary.

From the sensitivity viewpoint, the MCP-CCD, a charge-coupled device combined with a microchannel plate, may be one promising device for the detection of subtle image signals. Avalanche-multiplication type image tubes called HARPICON (Unnai et al. 1989) and super-HARPICON, developed recently for HDTV broadcasting, may also be interesting imaging devices. These have a wider dynamic range and an extremely high sensitivity, allowing a high-resolution color image at less than 1-lux illumination. If these were developed as solid state imaging devices in the future, there would be a possibility of revolutionizing machine vision applications.

Vision systems for complicated three-dimensional objects, objects with complicated surface patterns, and hard-to-see objects will also be needed in the future. The matching of distorted patterns (Sakou et al. 1988), a self-organizing image filter for texture separation (Sakou et al. 1990), and

a knowledge-based inspection of complicated patterns (Ejiri et al. 1989) are examples of recent approaches in these areas. Recognition of invisible objects is also becoming increasingly important. One such example is in the inspection of soldering conditions of surface-mounted LSIs on the circuit board, where the LSI leads are completely hidden under the LSI package. For such usage, visualization and image enhancement techniques are the keys for reliable vision systems. Obviously, the use of X-ray is one possibility for the visualization of hidden objects or hidden defects.

3.2 High-Speed Processing Techniques

Future applications will increasingly stress the importance of real-time-mode image processing, as the number of operations required to achieve a recognition purpose tends to be enormous. From recent trends in semiconductor technology, we can expect a 200-MIPS microprocessor (CISC-type) and a 1 gigabit memory chip in the late 1990s or early 2000s, if lithography limits and other difficulties in each generation are overcome smoothly. This means that the image memory with an adequate capacity is obtainable even for huge image sizes, which may eventually approach the human brain in size and capacity. This also indicates that the present supercomputer may be realized as a palm-top personal supercomputer. Therefore, optimistically speaking, almost all application problems today have a chance to be solved in the late 1990s or early 2000s.

However, it is usually a waste of time first to store the image into an image memory and then process it by accessing the image memory. Especially in industrial applications, it is a requisite to develop a fast method such that the completion of image acquisition implies the completion of the image processing (or at least preprocessing). Thus, the raster-scan type, one-pass real-time processing techniques based on on-site processing (Yoda et al. 1988, Yoda 1989) must be investigated further.

3.3 Sensor Fusion Techniques

In addition to vision systems, many other sensors must also be furnished, especially in a large-scale production system. The outputs from these sensors must be combined effectively to facilitate an optimal decision for intelligent manufacturing control. Thus, sensor fusion is becoming another key problem in the manufacturing process. The recent neural network approach may be one appropriate means to solve this sensor fusion problem.

A network simulator capable of simulating a network of millions of neurons has already been put

into laboratory use. In a hardware approach quite a few neurochips are being investigated in various institutions. One recent topic in this field is the prototype wafer-scale integration of a network with 576 neurons on a single wafer (Yasunaga et al. 1989). This indicates the possibility of integrating more than 10,000 neurons on a single wafer by the late 1990s. Neural networks in industry will be used primarily as vision devices to track the flow of parts and products and as sensor fusion devices for intelligent manufacturing control.

4 Other Applications

The 1990s is neither discontinuous with the 1980s nor with the 21st century. Thus, consideration of the 21st century will be important in determining what is necessary in the 1990s. Almost all problems facing the world today will become even worse in the coming century if steps are not taken to solve them now. These problems include environmental corruption, population explosion, natural resource shortage, and natural disasters. In some countries it is also important to cope with increasing social difficulties such as unfairness and uneasiness among people. Of course, these are political, economical, and social problems, and machine vision is obviously useless in solving them. However, there is a possibility to contribute somehow to the solution. Some application examples in this concern may be:

For clean and attractive environment:

- Environmental measurement and analysis
- Land and seabed cleaning
- Garbage treatment

For flourishing society:

- Food engineering such as agricultural automation, seabed cultivation, and automated animal breeding
- Efficient physical distribution control
- New traffic systems

For a safe and stable society:

- Traffic safety management
- Disaster prevention and rescue
- Security management
- Customs inspection automation

These applications may include the following new challenges in machine vision technology. They are the recognition of:

- Irregularly shaped objects and their defects
- Shape-changing objects and their degrees of deformation
- Quickly moving objects and their absolute/relative speeds
- Hidden/hard-to-see objects and the seriousness of their defects
- Flexible/soft/untouchable objects and their qualitative features expressed in such abstract words as matured, tender, colorful, lovely, and beautiful
- Individuals/living things and their existence, numbers, faces, and facial expressions
- Groups of living things and their static/dynamic behaviors.

The right directions necessary to solve the aforementioned problems must be decided, and clues to the solutions must at least be ascertained within the 1990s.

5 Conclusion

Futurology is always vague and involves many pros and cons in determining what is necessary for the future. Like other engineering problems, machine vision will also be expected to contribute somehow to the prosperity of future human society and the conservation of its natural environment. As described, there are few such contributions conceivable. Among them are the technologies to eliminate disasters, to increase human happiness by providing a comfortable living environment and a stable supply of essential commodities. We may have to think of these occasionally when conducting research in machine vision technology, and not just to think about publishing a paper.

References

- Ejiri M (1989a) Knowledge-based approaches to practical image processing. International Workshop on Industrial Applications of Machine Intelligence and Vision (MIV '89), Tokyo (April)
- Ejiri M (1989b) Machine vision: A practical technology for advanced image processing. Gordon & Breach Science Publishers, New York
- Ejiri M, Yoda H, Sakou H, Sakamoto Y (1989) Knowledge-based inspection for complex multilayered patterns. *Machine Vision and Applications* 2:155-166
- Sakou H, Matsushima H, Ejiri M (1990) Texture discrimination using self-organized multi-resolution filtering. *Transactions of IEICE, J73-D-II(4):562-573* (in Japanese)

- Sakou H, Miyatake T, Kashioka S, Ejiri M (1989) A position recognition algorithm for semiconductor alignment based on structural pattern matching. *IEEE Transactions of ASSP* 37(12):2148–2157
- Sakou H, Yoda H, Ejiri M (1988) An algorithm for matching distorted waveforms using a scale-based description. *IAPR Workshop on Computer Vision*, Tokyo (October)
- Unnai T, Goto N, Takasaki Y (1989) Newly developed high-sensitivity HDTV camera tube. *Hitachi Review*, 38(2):121–126
- Yasunaga M, Masuda N, Asai M, Yamada M, Masaki A, Hirai Y (1989) A wafer scale integration neural network utilizing completely digital circuits. In: *Proceedings of IJCNN*, vol. 25, pp 213–217 (June)
- Yoda H (1989) An algorithm for determining shape descriptors based on video-rate one-pass image processing techniques. *Transactions of IEICE, J72-D-II(7)*:1017–1028
- Yoda H, Ohuchi Y, Taniguchi Y, Ejiri M (1988) An automatic wafer inspection system using pipelined image processing techniques. *IEEE Transactions of PAMI* 10(1):4–16

ICPR Panel Discussion Remarks

Robert M. Haralick

Department of Electrical Engineering, University of Washington, Seattle, Washington

In order to design vision systems that work, a sound engineering methodology must be utilized. In the systems engineering approach a complex system is divided into simple subsystems, and from the input/output characteristics of each subsystem the input/output characteristics of the total system can be determined. Once an analysis methodology is firmly established, it becomes possible to work out a design synthesis methodology.

Machine vision systems are complex and are composed of different algorithms applied in sequence. Determination of the performance of a total machine vision system is possible if the performance of each of the subpieces, that is, the algorithms, is given. The problem, however, is that for most algorithms there is no performance characterization that has been established and published in the research literature.

What does performance characterization mean for an algorithm that might be used in a machine vision system? The algorithm is designed to accomplish a specific task. If the input data is perfect and has no noise and no random variation, the output produced by the algorithm ought also to be perfect. Otherwise there is something wrong with the algorithm. So, measuring how well an algorithm does on perfect input data is not interesting. Performance characterization has to do with establishing the relationship between the random variations and imperfections that the algorithm produces on the output data and the random variations and imperfections on the input data.

Now we are thrown into an immediate problem. It is typically the case that an algorithm changes the data unit. For example, an edge-linking process changes the data from a unit of pixel to a unit of a group of pixels. An arc segmentation and extraction process applied to the groups of pixels produced by an edge-linking process produces fitted curve segments. This unit change means that the representation used for the random variation of the input data set may have to be entirely different from the random variation used for the output data set. In our edge-linking–arc extraction example, the input data might be described by the false alarm/mis-detection characteristics produced by the preceding edge operation as well as the random variation in the position and rotation of the correctly detected edge pixels. The random variation in the output data from the extraction process, on the other hand, must be described in terms of fitting errors (random variation in the fitted coefficients) and segmentation errors. The representation of the segmentation errors must be natural and suitable for the input of the next process, which might be a model-matching process, for example. What should this representation be so that it becomes possible to characterize the identification accuracy of the model matching as a function of the input segmentation errors and fitting errors? Such questions have typically not been addressed in the machine vision literature. Until they are, analyzing the performance of a vision algorithm will be in the dark ages of an expensive experimental trial-and-error process.

Machine Vision in 1990s

Ramesh Jain

Artificial Intelligence Laboratory, The University of Michigan, Ann Arbor, Michigan

Machine vision has made significant progress in the last two decades. Many new techniques have emerged for many disparate applications. The most important change, however, has been in the availability of hardware related to all aspects of machine vision. In 1980 there were only a few places that could afford to have a laboratory with good image acquisition capability. Now anyone can buy image acquisition hardware and a computer to build a powerful vision workstation. This has resulted in research and development activities in many universities and industrial laboratories. It is expected that this trend will continue.

Machine vision has been applied to several problems. The following areas became popular and still are attracting lots of attention:

Inspection. Inspection has been a popular topic, and machine vision techniques have been applied from inspection of apples and oranges to solder joints on surface-mounted printed circuits. This area has still not seen as many vision systems as one would expect.

Medical imaging. Many aspects of medical imaging have attracted vision researchers, and now there are separate journals in this area. This is one of the faster maturing areas.

Robot vision. After a strong interest by the automobile industry in applying robot vision in its manufacturing, the popularity of vision systems in this area seems to be declining fast.

Navigation. Vision-based navigation has started attracting increasing attention in the last few years. This area continues to attract the Department of Defense as well as the automobile industry. Many other industries, including home robots, are showing an increasing interest in this area.

Document imaging. This may very well be the most important application of machine vision in the 1990s. There is tremendous enthusiasm in document processing. Many aspects of document processing are related to images. These images should be interpreted if document processing becomes really widespread. Many industries have realized this and are entering into this exciting area. Surprisingly few machine vision researchers have shown interest in research and develop-

ment related to this area. I expect this to be the most significant application of machine vision in the next and following decades.

Teleperception. Many applications require that a process be controlled from a remote location. In all these applications the work environment of a manipulator or a robot must be re-created at the location of the operator. This requires information assimilation using multiple sensors and then communication of this information. This application of machine vision has not received much attention yet, but increasing applications in space, environment restoration and waste management, ocean, and robotics in other hazardous environments are sure to result in many important developments in machine vision applications in the next decade.

An important question is how to facilitate application of machine vision to practical problems. I believe that this can be done by developing tools of *vision engineering*. Every applied branch has its scientists, technicians, and engineers. In machine vision we have only scientists and technicians. The role of engineers is to understand the strengths and limitations of different tools and techniques and to use them to solve a given problem. This usually requires that enough experiments will be done in controlled conditions to understand the scope of a particular tool. This tradition does not exist in machine vision. We need to develop an engineering culture in machine vision. Each operator should be tested under different conditions in order for us to have a good understanding of its characteristics. These characteristics will allow a designer to select a suitable operator for his or her application. Only then, we will not hear from our friend: "I am new to machine vision. For my application, I need an edge detector. There are so many of them! Which one should I use?"

I believe that applications of machine vision systems to disparate problems will result in the engineering culture. We have a strong research group in our field; now we need to pay attention to the development of vision engineering.

Finally, I believe that the 1990s will make machine vision a real technology. We will see the emergence of many applications, particularly in document imaging and image data bases.

Machine Vision in the '90s: A Technology View

Peter A. Ruetz

LSI LOGIC Corporation, Milpitas, California

1 Introduction

Machine vision problems can be divided into two primary classes: the front end and the back end. Of course, each of these groups may be further divided, but such fine distinction is not necessary. The front end is characterized by very regular pixel-oriented operations that must be computed at video rates (10+ Mpixels/sec). At the back end one finds more object-oriented operations that are much less regular and performed at a lower rate.

During the 1980s the emergence of a wide variety of real-time solutions to the front-end problems was witnessed. The most obvious problems were attacked, including convolution for filtering and edge detection, rank filtering for noise removal and morphology, histogramming and equalization, line detection, and FFTs/DCTs for analysis and compression, to name a few. In some sense the key term here is *obvious*, in that standard processors are built when there is sufficient agreement and consensus within the machine vision industry for the semiconductor industry to risk the investment associated with product development.

The back-end problems have fared quite differently in the same time period. There are few real-time integrated solutions available. A contour tracer is among one of the few object oriented devices now currently available (and even it is at the front of the back end).

2 What Happened?

The reason for the lack of standard products to perform the tasks required in the back end, in my mind, is that the machine vision community has failed to arrive at some kind of consensus as to what type of operations should be standardized. This standardization does not need to be formal; a de facto standard would suffice to bring out standard high-performance devices. Without these commercially available devices, each researcher must spend time redesigning software and hardware systems that could be standardized. In addition, the ability to process a large amount of data is hampered when all simulations are performed in software without real-time hardware. As a result, progress at the highest level of machine vision is retarded.

One may use the argument that any standard will be inferior to brilliant solutions of the individual researcher. However, any standardized solution should be at a much lower cost than the customized solution, and thus the standardized solution could include an array of operators, with the results being used in a custom way. In some sense one makes up in quantity for what one may lack in quality. At the same time the advantages of easy availability and low cost speed the remainder of the research.

3 One Example

I believe that a good parallel can be drawn from the video/image compression community. There are currently three major standards in the works: CCITT H.261, JPEG, and MPEG, each dealing with different video or image compression environment. In each case the experts in the field were able to set aside their differences to specify certain parts of the algorithm (e.g., each uses an 8×8 DCT followed by zig-zag run coding). However, they also left flexibility for various enhancements within a standard implementation. In other words, the "smartest" systems designers will end up with a better performing system that can still communicate with the poor performing systems.

The impact these standards had on the available hardware is striking. There are many companies (including LSI LOGIC Corporation) working to produce devices for these markets. With cheap devices available, the system designers can worry about finer details of the algorithm and not the computation of DCTs and motion vectors. We will likely see many systems based on these devices in the near future.

4 Conclusion

The machine vision community would be well served by some kind of standardization to promote the development of cheap high-performance processors. This would allow the designers to spend more time developing new capabilities and characterizing the performance of other systems and less time "reinventing the wheel."

Machine Vision Needs a Robust Technology

Jack Sklansky

University of California, Irvine, California

It is interesting to compare machine vision to computer graphics since they are complementary technologies: Machine vision automates the *analysis* of images, whereas computer graphics automates the *synthesis* of images. In this comparison it is clear that the technology of computer graphics has grown rapidly, whereas the technology of machine vision has grown much more slowly.

I suggest that the main reason for this disparity in growth is that *each machine vision system must be customized to a particular application* and that this customization is often expensive to carry out. For example, automating the analysis of fingerprints does not lead readily to automating the analysis of hand-printed characters, and automating the detection of defects in oranges does not lead readily to automating the detection of defects in microelectronic chips. On the other hand, customizations in computer graphics are relatively easy and inexpensive to implement.

I suggest, therefore, that the development of machine vision may be accelerated by facilitating and automating the customization process. This is equivalent to making the technology of machine vision more robust.

What do we mean by a “robust technology”? The word “robust” usually pertains to a *method* rather than a *technology*. We say a method is robust if it is effective in unanticipated applications. A technology consists of interrelated methods, devices, software tools, and, frequently, mathematical theories. We say a technology is robust if all of its components (methods, devices, software tools, theories) are also robust and if the interrelationships among these components are also robust (i.e., effective in the face of unanticipated applications).

Thus, a robust machine vision technology may include illumination invariant filters and noise-independent filters. In addition, adaptive mechanisms and processes may play important roles. For example, in segmenting an image, each pixel may be labeled by a classifier that adapts to each new application by an unsupervised training process (a form of cluster finding). Figures 1 and 2 illustrate such a training process. Figure 1 shows a design set of images consisting of dotted curves on a background of dotted noise. Figure 2 shows a test set in which the dotted curves are all dissimilar from those in the

design set and from each other. Figure 3 shows the results of an unsupervised adaptive classifier in which the dots in the dotted curves are joined by curves, whereas the dots in the dotted noise are not joined by curves. This illustrates the ability of adaptive classifiers to implement robust curve detectors. [Details of this form of training are described in Gutfinger et al. (1990) and Gutfinger and Sklansky (1990).]

A modern robust technology may also include computer-aided design tools. For example, the analysis of geometric relationships among objects in




Figure 1. A priori design images.

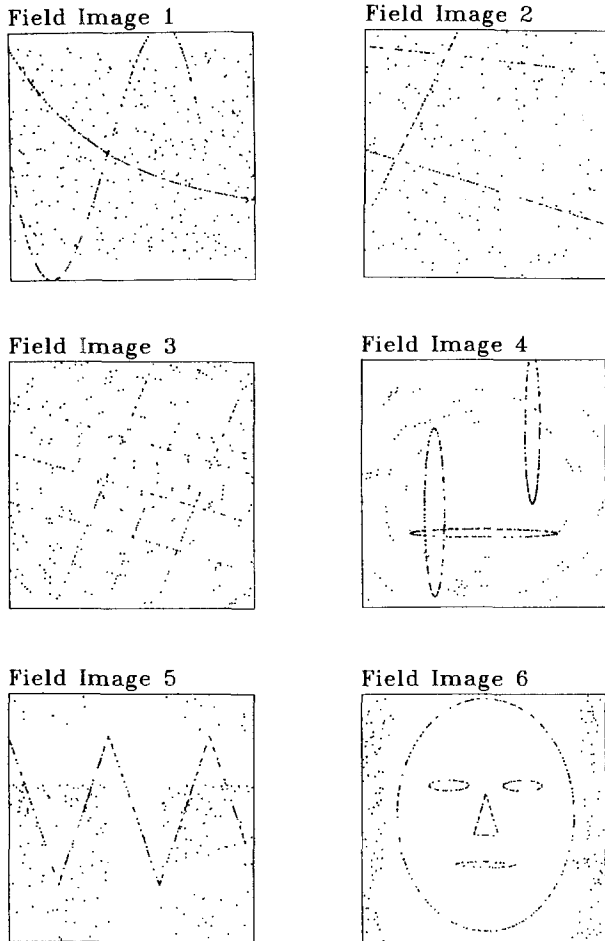


Figure 2. A posteriori field images.

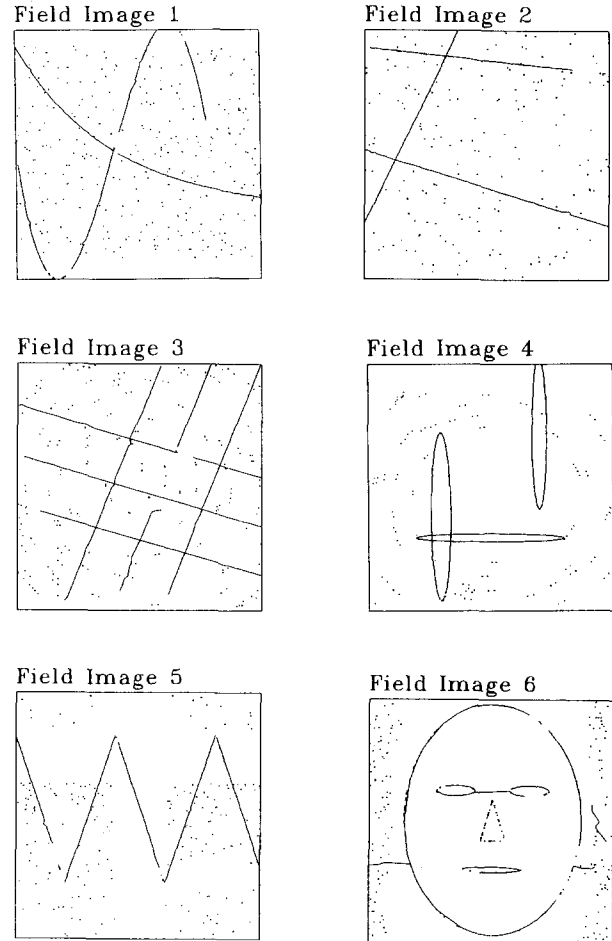


Figure 3. Results of application of adaptive classifier to curve detection.

an image may be facilitated by software for computer-aided modeling of these relations.

A robust technology may also include realistic mathematical models of transformations of scenes and noise that are likely to be encountered. Such models enable the construction of foreground analyses algorithms that will perform well in a wide range of noisy backgrounds. Edge-preserving noise suppressors are examples of algorithms that can be made robust by exploiting models of edges and models of noise.

Another feature of a robust technology is that the components of the technology are well integrated regardless of the application.

In summary, a robust technology of machine vision may include components such as the following:

1. Illumination-independent filters
2. Noise-independent filters
3. Adaptive pixel classifiers for low-level processors—for example, edge detectors

4. Computer-aided design of analyzers of geometric relations among detected objects
5. Adaptive classifiers of intermediate objects—for example, curves and regions
6. Three-dimensional modeling of objects appearing in multiple views of the same scene.

The effective integration of these components, we believe, will give us a form of systems engineering that will suppress the ad hoc nature of machine vision techniques, and thereby accelerate the growth of modern machine vision technology.

References

Gutfinger D, Nishimura R, Doi H, Sklansky J (1990) Robust curve detection by temporal geodesics. Technical Report TP-90-3, Pattern Recognition Project, Depart-

ment of Electrical and Computer Engineering, University of California, Irvine, CA, (April) also, in: Proceedings of ICCV'90, Third International Conference on Computer Vision, scheduled for December 4-7, 1990

Gutfinger D, Sklansky J (1990) Robust classifiers by mixed adaptation. Technical Report TP-90-2, Pattern Recognition Project, Department of Electrical and Computer Engineering, University of California, Irvine, CA (April)

Machine Vision Trends Entering the '90s

C. W. Swonger

Image Processing Systems Division, Environmental Research Institute of Michigan (ERIM), Ann Arbor

The 1980s can be viewed with substantial supporting evidence as the decade in which industrial/commercial machine vision was subjected to trail by fire. As a national technological business sector, machine vision began the 1980s with an enthusiastic rebirth (after false starts in the 1960s), commanding the attention of technologists, venture capitalists and potential users worldwide. In the mid-1980s literally hundreds of machine vision product companies appeared and disappeared, as technology developments failed to find receptive markets able to assimilate and/or cost justify and/or afford machine vision products. The late 1980s have seen significant consolidation of the industry; emergence of selective, cautious, and better "educated" users; and some signs of sounder technical and applications engineering approaches to the industrial and commercial applications of machine vision.

Predictions of the future of any dynamic technology are always difficult and risky. Nevertheless, it may be worthwhile to examine current machine vision trends so that plans and decisions can be made that preserve future options and maximize the probability of beneficially anticipating future developments. To a considerable degree, future industrial vision trends can be detected by examining the continuing rapid advancements in machine vision technology arising from *government* application of machine vision (i.e., image processing, remote sensing, and target recognition). From this arena of aerospace/defense technology many past advances in industrial automation capability have been spawned.

What are some of these trends? They can be grouped into sensors, software, processors, algorithms, and applications.

Sensors

- Advanced focal plane sensor arrays have appeared. These will continue to extend the spatial

resolution, the radiometric sensitivity, and the dynamic range of vision sensors that use such devices.

- Advanced high-information sensors will continue to achieve new levels of performance and lower costs as the user community learns how to exploit them in high volumes. A prime example of such sensors is the generic class of laser radars (LADAR) that, in various forms, can provide direct three-dimensional measurements, color discrimination, and motion measurement.

Software

By software we refer to the system control, operating systems, and solution development tool programs that are used in machine vision systems, not the algorithms that specifically analyze the sensor scene or image. These often overlooked, but vital, elements of vision systems have frequently been the downfall of past vision company offerings. High "nonrecurring engineering" costs experienced by machine developers can usually be traced to inadequate development environments, testing tools, or system software facilities to allow easy reconfiguration or interfacing to manufacturing systems. Trends include:

- Emergence of standard and more general very high-level languages that dramatically enhance developer productivity and portability of algorithms to ever-advancing hardware suites
- Maturing of generalized user and display environments permitting iconic or "visual" programming and output data presentation
- Integration of sophisticated imaging simulation systems that allow algorithms to be rapidly and knowledgeably stressed during their development so as to probe their performance and ensure their robust behavior

Processors

- Processor architectures and devices have also evolved to reflect the value of a requirements-driven systems engineering approach to a machine vision problem solution.
- The inadequate general purpose processors and the powerful but highly stylized homogeneous special purpose processors of the 1980s are being supplanted by heterogeneous vision processor systems in the 1990s. These reflect the inherent character of the generic vision application. That is, high-intensity processing of moderate precision pixels, rule-based artificial intelligence paradigms, and high-precision matrix/vector processing are each accommodated efficiently and with proper balance, and each in specialized engines within a loosely coupled multivision processor system.
- Electronic packaging technology is beginning to emerge that will make feasible the goal of compact "brilliant cameras." The size and cost of such units will permit their insertion at many observation points in manufacturing processes. These will replace the cabinets and boxes that have required excessive power, space, and maintenance.
- Together, the preceding heterogeneous architectures and densely miniaturized packaging are already making feasible two to four decimal orders of magnitude increase in processing capacity per unit volume in aerospace/defense vision processors. This will translate into more robust, reliable, and accurate industrial vision performance over the next decade.

Algorithms

Algorithms can be viewed as the heart and soul of a machine vision system where adequate measurements or discrimination or control either is or is not achieved.

- While still largely an art, machine vision algorithms and their development methodologies are beginning to incorporate greater use of systems engineering principles and mathematical formalisms, such as the generalized image algebra of Ritter.
- After years of diverse piecemeal approaches to image enhancement, detection, measurement, and recognition (both in industrial and aerospace/defense applications), some convergence and integration of algorithmic paradigms is occurring. Thus, more developers of solutions are coming to

realize the benefits of combining statistical, neural network, morphological, artificial intelligence, and signal processing algorithm components to create more effective complete vision systems. Unfortunately, there still exist various camps that proclaim the universal superiority of their special notion or technique.

Applications

- Successful applications of machine vision, in a technological sense, steadily grew through the 1980s. However, successful applications in the sense of those on which viable high-volume machine vision business could be based were "few and far between." In addition to various other socioeconomic factors, this was largely due to the excessive cost of so-called "nonrecurring engineering" (which often turned out to be, in fact, terribly recurring). This was in turn due to the fragility of the ad hoc algorithms, sensing/illumination, and system architectures employed in machine vision products. Thus, only the few very uniform high-volume applications that could be identified yielded some semblance of a repeatably salable product line.

To the extent that present and future machine vision business ventures mature away from these costly limitations, the 1990s should see a significant increase in reliable businesses based on machine vision applications. The greatest needs and opportunities probably exist in:

- Consumable products analysis, inspection, and processing, such as:
 - Pharmaceuticals
 - Food products
- Electronic components and assemblies inspection, fabrication control, assembly, and testing, such as:
 - Chip packages
 - Multichip substrates and assemblies
- Durable goods manufacturing:
 - Fit and finish inspection and process control
 - Vision-guided robotic assembly
 - Calibration and alignment process control
- Automated document processing:
 - Handwritten document data
 - Integrated processing of text, sensor-derived imagery, graphics, and other data

Finally, any such application, to be successfully exploited, will continue to require absolutely the dedicated collaboration and commitment of:

- Manufacturing and commercial end-users who are expert in their processes and products and also educated and receptive to the technical trade-offs in machine vision
- Machine vision suppliers who have the persistence and insight to understand and adapt to the real requirements of the manufacturing or commercial user, and also to bring the more flexible

and powerful machine vision technology and tools described earlier to bear on the user's requirements.

As this bilateral integration of disciplines occurs, machine vision will finally begin to realize its full potential on a national and worldwide scale.