Industrial machine vision: lessons and challenges

By Steven W. Holland and Robert B. Tilove

A review of critical technologies that have contributed to the success of industrial machine vision applications and potentially significant trends.

Although research in machine vision began in the late 1960s, industrial applications have only recently emerged. Most forecasts regarding the future growth of commercial machine vision are highly optimistic. We believe it is both instructive and timely to reflect on the state of the art in an effort to understand, on the one hand, the key technical and economic factors that contribute to successful industrial applications, and to identify, on the other hand, important lessons and/or technological barriers that should influence current research and development trends.

General Motors has been active in machine vision for 15 years: work began in the Research Laboratories (GMR) in the early 1970s. In addition to its ongoing research effort, GM also has today a complementary activity within its Advanced Engineering Staff, which is seeking to accelerate the implementation of machine vision technology throughout the corporation through a combination of in-house development and strategic relationships with outside vendors.

With this rich base of experience, it would have been easy to focus solely on GM applications of machine vision in preparing this report, but we sought a more general view. We contacted a number of developers of sophisticated machine vision systems and requested their input on systems they regarded as exemplary of the state of the art. From this information and our internal experiences, we attempted to extract some common threads and unifying themes.

We found that systems widely regarded as "landmark" or "state of the art" shared many of the following key characteristics:

- They solve "just the right set" of problems.
- They utilize appropriate, sometimes novel, computational methods.
- They utilize appropriate, sometimes novel, technological tools.
- They incorporate sound basic system design and engineering practices.

In what follows, we shall elaborate on each point with references to information provided by our contacts. The particular systems we elected to use as examples were chosen mainly because they seemed appropriate for emphasizing the main points of interest. This report is not a survey of industrial applications, as the number of impressive industrial vision systems today is simply too large to attempt a comprehensive study. We hope that our readers will be forgiving if they find that we have neglected to mention their favorite system or if our list of key characteristics does not entirely match their own.

Application scope

A striking feature of successful vision systems is that they solve just the

STEVEN W. HOLLAND is an assistant department head in the Computer Science Dept. of General Motors Research Laboratories (GMR). He is responsible for research projects involving computer vision, robotics, mobile robotics, process planning, knowledge systems, natural language, advanced programming environments, parallel processing, and statistics.

ROBERT B. TILOVE is senior staff research scientist and project leader of the machine perception project at GMR. His research interests include geometric modeling, model-based programming, and control of intelligent automation systems.

right set of problems. The application scope of a vision system reflects a trade-off between generality on the one hand and ease of use and robustness on the other.

Most potential users of machine vision would like some sort of general purpose vision system that can easily be configured to perform a wide variety of tasks. But this is not possible for both technical and economic reasons: the "general vision problem" is not well understood, and even if it were, general purpose tools would probably not be as effective as tools specifically designed for a small class of similar applications.

Current applications of machine vision tend to attain reliability, robustness, ease of use, and maintainability at the expense of generality . . . The result is specialized tools that are like grapefruit knives: they solve one problem very well but are not easily applied to new or unanticipated problems.

Current applications of machine vision tend to attain reliability, robustness, ease of use, and maintainability at the expense of generality. As a result, the application scope of most successful machine vision systems is quite small. The result is specialized tools that are like grapefruit knives: they solve one problem very well, but are not easily applied to new or unanticipated problems. We shall describe below three examples of systems that were carefully designed for particular problems.

GM's KEYSIGHT system,[1] illustrated in Fig. 1, is a gray-level system that inspects engine valve spring assemblies for the presence of assembly retaining keys. It first extracts edges, a reasonably generic step, but thereafter, KEYSIGHT's operation is application specific. A symmetric image is assumed, but the center of symmetry may not be known precisely and is therefore computed using autocorrelation methods. Having found the center of symmetry, KEYSIGHT integrates image intensity (in the original gray-level image) along a circular contour whose radius coincides with the center of the expected retainer keys. Images of assemblies with missing keys can be reliably distinguished

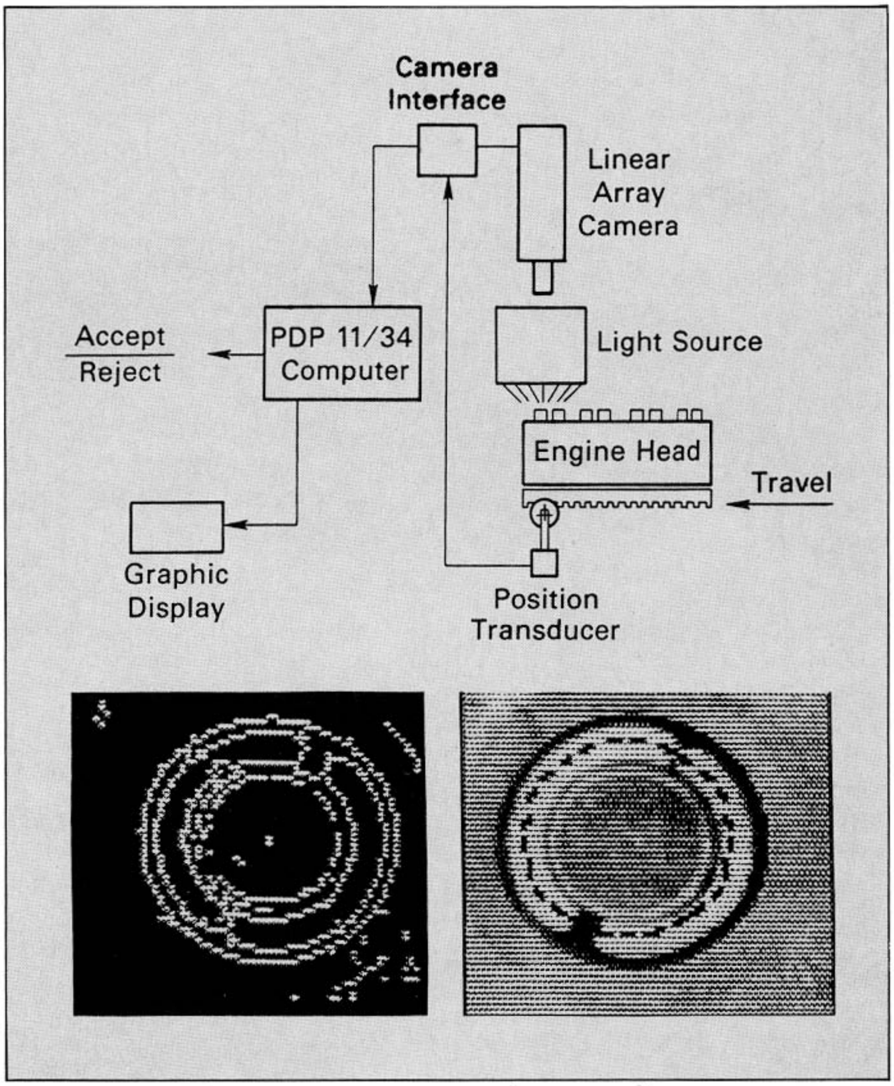

FIGURE 1. GM's KEYSIGHT system tests for the presence of a particular component in a particular assembly. The system architecture is shown in the upper portion. The lower portion shows two of the processing steps: edge detection (left) and inspection based on the integral of image intensity along a circular path (right).

from images of good assemblies on this basis. Because GM has a number of similar engine assembly lines, KEYSIGHT enjoys a large potential market.

Robotic Vision Systems Inc.'s (RVSI) APOMS (Automated Propeller Optical Measurement System) is also tailor-made for a particular application—but in this case, the potential market is quite small. APOMS was developed specifically for the Naval Sea Systems Command to measure 24-ft diameter, 50-ton submarine propeller blades with high accuracy over large spaces.[2] (It measures 100 points per square inch to .0025-in. accuracy, and requires 10 hours per blade.) A special robot was designed specifically to carry the sensor head (Fig. 2). Two other interesting aspects of APOMS are: a) the inspection points are taught off-line from a CAD-style description of the blade surface; and b) APOMS will eventually be integrated with a grinding machine to correct inaccuracies.

Vision systems that currently enjoy the largest customer base are found in the electronics industry. The systems have been carefully designed to meet specific needs in the industry, are easy to use and program, and interface comfortably with existing automation. In this sense, these systems exemplify careful selection of application scope.

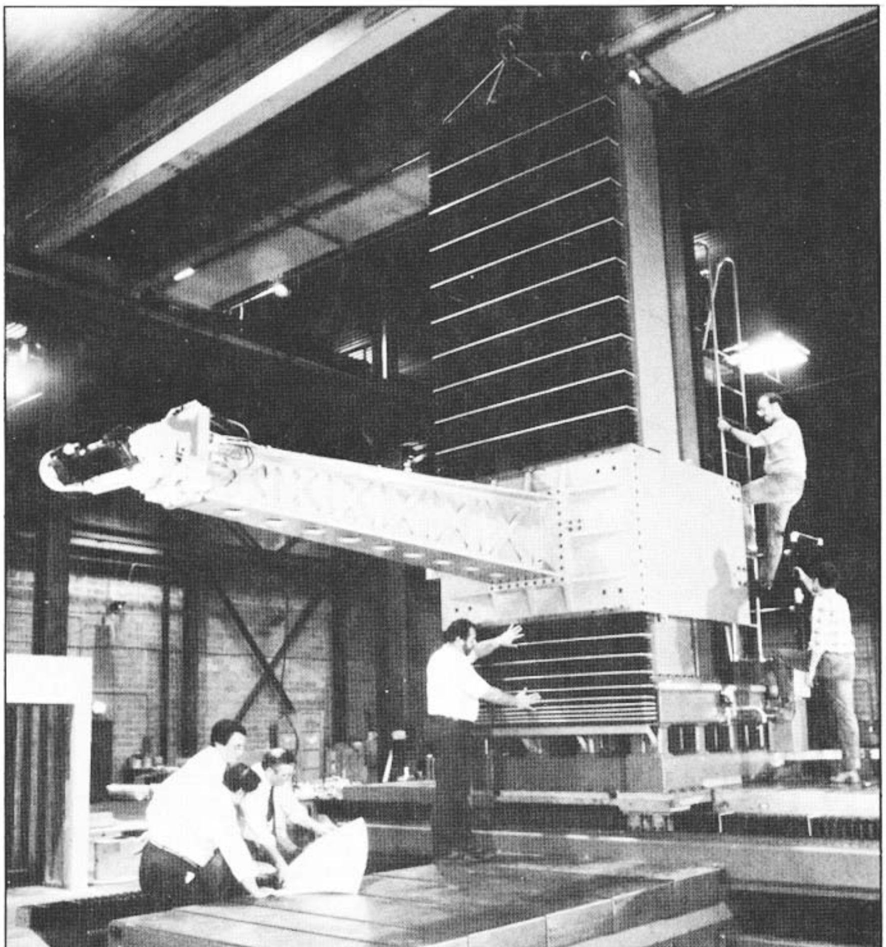

FIGURE 2. The APOMS system includes a custom built robot for scanning the sensor (camera and light source) over the propeller surface.

That such systems are found mostly in the electronics industry may be due to the fact that the industry is highly automated, competitive, and capital intensive, and also because the problem is simplified somewhat by clean environments and two-dimensional inspection problems. View Engineering, Contrex, and KLA are examples of vendors specializing in such applications. Baird has provided a good overview of the applications and approaches.[3]

Computational methods

Application scope defines a class of problems to be solved; it does not define the computational methods employed in the solution. The primary goal in computer vision is to extract automatically, from physical sensor data, intrinsic properties of the objects in the underlying scene. But general "image understanding" is well beyond our ability to define, let alone solve. In practical applications, vision systems exploit constraints, both natural (intrinsic to the underlying objects) and artificial (imposed externally by the imaging system), and expectations to interpret the data.

By natural constraints we mean those properties of an image that rely on reflectance, color, texture, and other physical properties of the specific objects to be analyzed. Natural constraints can often be exploited in the method of inspection or measurement. The ERIM (Environmental Research Institute of Michigan) thick-film inspection system, for example, utilizes specific reflectance properties of the film to establish inspection criteria.[4]

By artificial constraints we mean those properties of the image that rely on the vision engineer's choice of lighting, lenses, and so forth. A common example of artificial constraint exploitation is "structured light." By illuminating an object with a bright line of light, image intensity data can be converted into range data, and easy-to-locate image features can be produced where there otherwise would be none.

In GM's CONSIGHT system, structured lighting is used to produce binary range images of objects on a moving conveyor belt (see Fig. 3).[5] In GE's turbine blade inspection system, a scanning structured light system is used to produce profile measurements.[6] Vendors such as Diffracto, Perceptron, and RVSI specialize in structured light applications. Colored (rather than geometrically patterned) lighting has also been used, for example by GE[6] and Sony.[7]

Other types of expectations obviously play a critical role in machine vision: measured data is compared to expected values or is fit to predefined or learned models, particular brightness patterns (templates) or features are located, and so on. We are beginning to see industrial vision systems in which predictions are derived from a CAD database describing the ideal shape of objects in the scene, e.g., RVSI's APOMS system, previously mentioned.

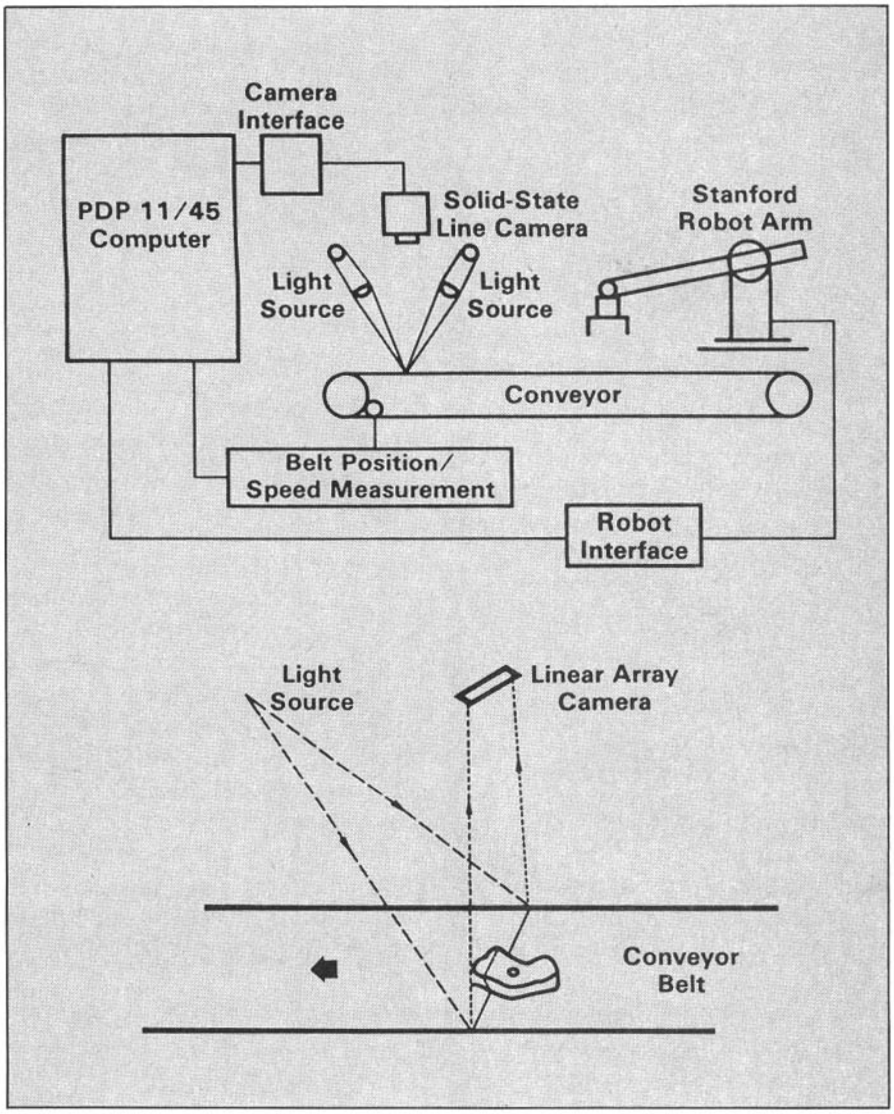

FIGURE 3. CONSIGHT uses a structured line-of-light imaging system to recognize and locate objects on a moving conveyor belt. The overall system architecture and the lighting system concept are shown.

The use of constraints and expectations are what enable practical machine vision systems to operate satisfactorily and are also one of the reasons that the application scope of practical systems is small. Much of the application-specific engineering and creativity that underly industrial applications seems bound up in exploiting or imposing the right set of constraints. Indeed, there have been cases where parts were redesigned and/or manufacturing processes revised for the sole purpose of facilitating subsequent visual analysis.[8]

Technological tools

Technological tools relate to the implementation of vision systems and include digital image processing hardware, optical image processing hardware, electro-optical transducers, and light sources.

Advances are being made in the use of custom digital hardware for efficient binary and gray-level image processing.[9-12] The main driving force for this is that digital image processing operations (e.g., filtering and correlation) are local operations that are computed over a large image. Software implementations on general purpose processors are often too slow for practical applications. The operations lend themselves to implementation in custom parallel/pipelined hardware. The View Engineering system mentioned earlier uses a custom processor to correlate live images with pre-taught templates, as do GMR's Seamsight[13] and Multi-Match[14] systems.

High-density, low-noise and distortion, electro-optical transducers are becoming available, and these have paved the way for more accurate, higher resolution vision systems.

Optical image processing techniques have produced dramatic results for surface defect measurement in systems like GM's FLAWS[15] and Diffracto's DiffractoSight.[16] In FLAWS, for example, the image of a coherent line of light is optically filtered to remove low-frequency components. High-intensity portions in the result are indicative of surface irregularities.

Good basic design practices

Industrial grade systems obviously must be designed according to sound engineering principles. They must be reliable; robust in the presence of noise and other degradation of data; and easy to calibrate, operate, and maintain.

Early systems were forced to rely on cameras, computers, and other hardware that were not intended to withstand the factory environment. Today we are beginning to see vendors such as Perceptron[17] package sealed units containing light sources and sensors that have been pre-calibrated and optically aligned. External precision mounts enable sensor units to be replaced with minimal system downtime. Computers and communication hardware are enclosed and designed to withstand harsh environments.

Initial laboratory-style systems clearly demonstrated the feasibility of vision technology, but rigid design standards are obviously needed for production-hardened systems. In typical industrial applications, one finds tons of concrete supporting cameras and fixtures, armored conduit, and NEMA-12 enclosures housing custom computers and AC line filters. To vision scientists, these details are, perhaps, uninteresting. To manufacturing engineers, and to vendors who wish to sell systems, these issues are paramount.

Revolutions on the horizon

Research in motion, stereo, image smoothing, edge and line detection, neighborhood processing, and other computational approaches to vision will eventually lead to vision systems that are far more robust, easy to use, and less dependent on artificial constraints than current systems.

To date we have made only limited use of computer-based object design data in our vision systems. The integration of CAD data with vision will untimately lead to systems that are easy to program and self-calibrating.

Vision is inherently parallel, and the eventual availability of general purpose highly parallel computers will be a fundamental breakthrough, not simply because of dramatic speed improvements, but also because entirely new computational approaches will become practical.

Finally, new kinds of sensors, for example, those that produce accurate range data at speeds and spatial resolutions comparable to today's solid-state TV cameras, will provide new sensing modalities that do not suffer from some of the limitations of traditional TV cameras.

Lessons

Among the lessons to be drawn are:

- It is not a cheat to do good engineering. In the same way that a professional chooses the right tool for a job, vision engineers must exploit constraints and expectations to achieve the robust, reliable performance demanded for industrial-grade systems. The first-generation systems in our plants today represent significant technical achievements that evolved over a relatively short period of time. We have every reason to be optimistic about the future.
- Because of the magnitude of the engineering and development efforts involved in the design and implementation of a system, the application scope must be carefully chosen to permit multiple installations. We view the general lack of vision building blocks as evidence of a lack of maturity in the field, but increased pressures to automate and to improve quality, combined with growing numbers of installed applications, are likely to encourage modularity.
- To date there has been little direct interaction between sensor developers and vision system developers. Intelligent sensors of the future should produce pre-processed data (calibrated, scaled, and distortion- and quirk-free) in engineering units appropriate to the application. Development of such sensors will obviously require close collaboration.
- Vision research should focus on technologies likely to revolutionize industrial vision: computational theories of vision, CAD interfaces, range and other new sensors, and parallel hardware.

Acknowledgments

We gratefully acknowledge information provided by: Toshio Konishi of Sony Corp.; Russell Petersen of Synthetic Vision Systems; Michael Baird of IMC Corp.; James West of Perceptron; Ray Davis of Chesebrough-Pond's Inc.; Howard Stern of Robotic Vision Systems Inc.; Yoshiaki Shirai of Electrotechnical Laboratory; Masatsugu Kidode of Toshiba Corp.;

Harlan Neuville of General Motors; Tim Pryor of Diffracto Ltd.; and Peter Wohlmut of KLA Instruments Corp.

REFERENCES

1. N.A. Perkins, D.B. Juliette, R.M. Iadipaolo, J.K. West, "KEYSIGHT: A computer vision system for inspecting valve spring assemblies," in *ASME Automation in Manufacturing Systems Processes and Computer Aids* (P.E.D. Vol. 4, 1981).
2. H. Stern and R. Metzger, "Automation of propeller inspection and finishing," Nav. Eng. J. 97 (4) 124 (1985).
3. M. Baird, "EYESEE: A computer vision system for inspection of integrated circuits," Proc. Soc. Photo-Opt. Instrum. Eng. 470, p. 261 (1984).
4. D.J. Svetkoff, J.B. Candlish, and P.W. Vanatta, "High resolution imaging for automatic inspection of multi-layer thick film circuits," paper MS83–222, presented at the SME Applied Machine Vision Conference, Memphis, Tenn., 1983.
5. S.W. Holland, L. Rossol, and M.R. Ward, "CONSIGHT-1: A vision-controlled robot system for parts on belt conveyors," in *Computer Vision and Sensor-Based Robots*, G.G. Dodd and L. Rossol, eds. (Plenum, New York, 1979), p. 81.
6. A. Porter and J. Mundy, "A non-contact profile sensing system for visual inspection," paper presented at the Workshop on Industrial Applications of Machine Vision, (IEEE Computer Society, Research Triangle Park, N.C.), p. 119 1982.
7. M. Takagi, T. Konishi, and J. Kitsuki, "Shape reconstruction of wires using colour images for automatic soldering system," presented at the International Symposium on Robotics Research, Bretton Woods, N.H., 1983.
8. H. LaCoe, "Working together on design for vision," Vision **1** (1), 7 (1984).
9. E.W. Kent, M.O. Shneier, R. Lumia, "PIPE (Pipelined Image Processing Engine," J. Parallel and Dist. Computing, 2, 50 (1985).
10. Synthetic Vision Systems, Ann Arbor, Mich., product literature.
11. M. Kidode, "Image processing machines in Japan," Computer **16** (1), 68 (1983).
12. View Engineering, Simi Valley, Calif., product literature.
13. I. Masaki, "SEAMSIGHT: A parallel-pipelined vision system for seam tracking," in *Proceedings, Seventh International Joint Conference on Pattern Recognition* (IEEE Computer Society Press, New York, 1984), p. 424.
14. I. Masaki, "Parallel/pipelined processor dedicated to visual recognition," in *Proceedings, International Conference on Robotics and Automation* (IEEE Computer Society Press, New York, 1985), p. 100.
15. A. Gara, "Automatic microcircuit and wafer inspection," Electronics Test **4** (5), 60 (1981).
16. Diffracto Ltd., Windsor, Ontario, product literature.
17. Perceptron Inc., Farmington Hills, Mich., product literature.