The Promise and Payoff of 2D and 3D Machine Vision: Where are we today

Kevin Harding
General Electric Global Research Center
PO Box 8, Schenectady, NY 12301

ABSTRACT

In the past 25 plus years, machine vision has grown from a high priced solution looking for a problem to solve, to a multi-billion dollar industry playing a crucial role in todayís high demand production environment. Early machine vision systems consisted of dedicated, special architecture processors, some with hundreds of individual processors, and price tags approaching a hundred thousand dollars or more. There were large array boxes that required high end workstations to communicate with, dedicated units with control like interfaces the size of small refrigerators, and special low level languages understood only by the most dedicated programmer. Today, a full featured vision processor will fit into a standard PC box, often as a plug in card, and the fast, dedicated purpose systems will fit in the palm of your hand, connecting to any PC over an internet connection. The 3D vision system has likewise made great strides, though it remains just a step or two behind its 2D cousin. Early 3D systems were notoriously slow, taking the good part of an hour on temper mental equipment, producing complicated clouds of data with not good way to use the information. Today, high quality 3D data can be obtained from rugged, even portable units with simple to use interfaces, producing inspection information based upon part tolerances and CAD models in a matter of seconds. None of these systems are the mystical ìRobots that can seeî envisioned in those early days. But what machine vision is today, and what it is becoming, is a technological tool on its way to becoming as common place as the computer, not only in production environments, but potentially in our every day lives. This paper will look at what 2D and 3D vision can do today, where and how it is being used, and where it may be going in the future.

Keywords: structured light, three-dimensional, machine vision

1. INTRODUCTION

Modern day production lines are making and moving parts at speeds much faster than any other time in history. The standards of six sigma quality have demanded much better control than ever over even small, cosmetic defects. Industries such as primary metals, automotive, textiles, even plastic extruders have found that having about the right dimensions and being ëfunctionalí just isnít enough. Manufacturers are finding that any appearance of quality problems, be it pits and scratches or a bad overall appearance can mean rejections of full lots of product, costing millions of dollars to a company, and affecting their bottom line. In the most critical industries, over 80 percent of customer rejects are often due to non-function impacting, cosmetic defects.[1] At modern speeds of production and tight defect tolerances, human inspectors have trouble keeping up to production. Studies have shown that even after 2 hours of such work, the human inspector becomes distracted. The same mind that provides for high defect discrimination, can "fill in" missing pieces, even when they are not present. After seeing 1000 parts with a hole in the center, part 1001 will appear as though it has a hole, whether it does or not.

Computers and the internet have provided the tools to deal with large amounts of information very quickly. The same limitation that requires that a task be completely spelled out for a computer, insures that it will find that missing hole in part 1001 as consistently as in part 1,000,001. In addition, the simple act of reporting a variation, inherent to the philosophy of statistical process control, becomes a quick transfer of data over internet lines, in the digital form needed for SPC software. So, computer based inspection and monitoring, not only affords the programmable flexibility demanded by flexible manufacturing, but also provides the quick data collecting and tracking abilities needed for high speed repetitive operations.

Two- and Three-Dimensional Vision Systems for Inspection, Control, and Metrology,
edited by Bruce G. Batchelor, Heinz Hügli, Proceedings of SPIE Vol. 5265 (SPIE,
Bellingham, WA, 2004) · 0277-786X/04/$15 · doi: 10.1117/12.523472

But why use optical based inspection and gaging? Manufacturing has employed contact probes and gages in regular use since the turn of the 20th century. Coordinate measurement machines (CMMs) have gone from slow, laboratory systems to automated factory floor systems. But even with those improvements, 100 percent inspection is rarely feasible with CMMs alone. Many fixed gages have now become computerized as well, providing a dedicated part gage, with computer output at the speeds needed. For loading these gages, robotic systems are able to load and unload parts in a highly repeatable manner, so good that they have revolutionized the electronics fabrication industry. But this option means a dedicated set of gages for each part, demanding rooms full of gages, and billions of dollars in expenses each year. At these costs, the small batch run envisioned as the main tool of flexible manufacturing systems just is not economically feasible. Even with these computerized advances, the high speed and high tolerances of new parts has pushed past the limits of these more traditional sensors. The combination of speed and resolution has hit mechanical limits that even with new light weight materials have not yet been overcome. The flexibility of machine vision to check hundreds of points on one part, then a different set of points on the next part, all in a matter of seconds has provided a capability not before available with traditional fixed gages.[2]

The progression of machine vision as a tool in manufacturing has not been an overnight occurrence. Early applications of machine vision as a sorting tool and part ID aid were little more than hundred thousand dollar bar code scanners. High speed, low cost, and flexible change over in the fast moving computer and semi-conductor industries has acted as a catalyst to increase the speed of these machine vision system as well as increasing the sophistication from sorters to high precision measurement tools working in the micron regime. Early machine vision systems using simple processor chips progressed to dedicated ICs, gate arrays, digital signal processing (DSP) chips and now integrated internet devices. The dynamic nature of the electronics and semi-conductor market segment has kept these areas as the largest current application of machine vision, still accounting for over 50 percent of sales in a multi-billion dollar worldwide machine vision market today. Even 30 years after the early factory applications of machine vision, many durable goods manufacturing plants have not embraced the technology. But the competition for tighter quality control may push vision technology in to even the most conservative metal cutting and forming operations.

There have been quite a few other applications that have proven themselves in the industry, as well as those that have not panned out. The checking of packaging in critical industries, particularly pharmaceuticals, has proven an important tool in satisfying government guidelines. Checking assemblies such as bearings for missing parts has saved this industry millions in failures. Some of the tougher problems such as finding cosmetic flaws in patterned textiles, metals or plastics have proven more elusive, requiring a degree of reasoning not as easily provided by a computer. The high speeds of computers today coupled with extreme overseas competition has been pushing vision into even these areas. The alternative has proven to be loss of business or even bankruptcy for those who have not made the transition.

Table 1. Application areas of Machine Vision[3]

APPLICATION	STATUS	REL. VOLUME
Electronics/semiconductor	still fast growing	~ 52 %
Packaging	Pharm. growing	~ 10 %
Character Reading	standard tool	~ 14 %
Food and Agriculture	Stable market	~ 12 %
3D Gaging	next big wave	~ 11%

So what is this machine vision and what is needed for a successful application?[2] Machine vision can actually be any system where visual information is sensed and analyzed by a computer system for the purpose of determining something about the manufacturing process[4] Such a system typically consists of a light source to highlight the features, some optics to image, a video camera to sense the scene, a digitizer to move the video into the computer format, and a computer system with analysis software (see diagram in Figure 1).[5-7] By industry consensus, machine vision strictly relates only to the application of this technology to the manufacturing environment, which has been the mainstay of vision technology for the past 25 plus years. However, it is very telling of the maturity of the technology that applications in other areas such as medical image analysis, transportation, and security are being seriously pursued today.

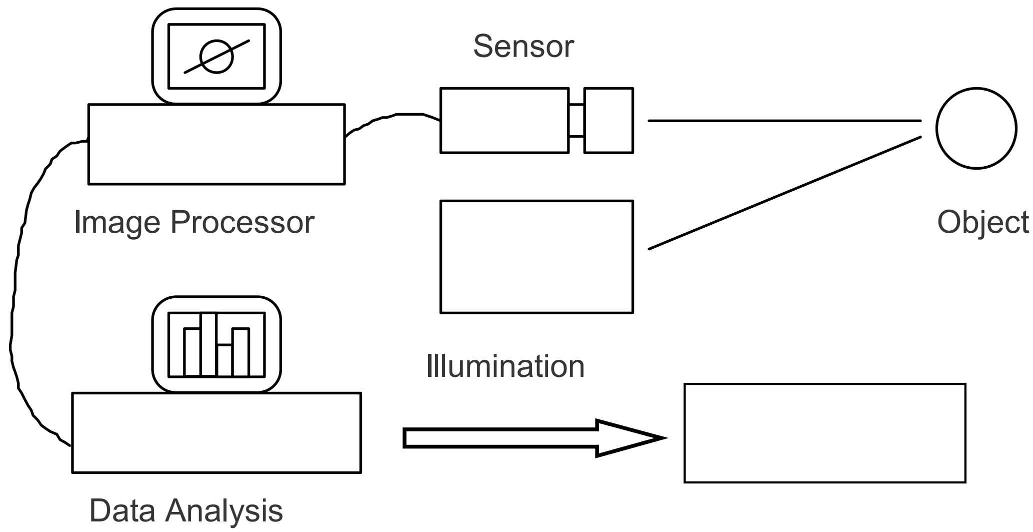

Figure 1. Block Diagram of a typical Machine Vision System.

The lighting and optics system creates an image on the video camera of some Region of Interest (ROI) or Field-of-view (FOV). The image is recorded by the video camera and digitizer to create an array of picture elements or "pixels", each of which denotes one point in the image. The image might then be smoothed by means of averaging or filtering, then segmented to isolate key "features" of interest.

The value of each pixel may be a grey level of light, or may be binarized to be either dark or light. The analysis then typically uses simple tools that relate the pixels to each other by identifying edges, counting the number of dark or light pixels in an area, or looking for certain patterns in the arrangement and values of the pixels.

There has been many advances made in dedicated computer systems and user friendly software over the years. In fact, machine vision owes a great deal of its current success to the explosion of the personal computer industry. The large consumer market for computers has brought the price tag for machine vision systems down by an order of magnitude in the past ten years. Easy to program icon based systems for computers have also helped to provide much easier to use software for machine vision. Some of the new vision algorithms have just been refinements of old ideas on new, faster computer platforms, where others are new ways of doing things. Camera systems have also improved over recent years, fueled by the use of home digital cameras. The combination of camera technology and computers has facilitated a whole new area of machine vision referred to as smart cameras. These smart cameras have on-board processors which can perform a powerful, if limited set of operations for such defined operations as part ID, location, and simple measurements, then communicate the results over an internet connection. Let us consider where the technology of machine vision is today, and where that technology may be taking the field in the future.

2. VISION TECHNOLOGY DISCUSSION

2.1 Lighting and Optics

The area of machine vision has seen quite a few advances in the past 25 years. The first step, and perhaps the least developed today, is getting a good image to analyze. The problem is not that the technology for producing high quality imaging does not exist, it is more that the machine vision market has not as yet reached the type of commodity volumes of say home cameras that make investing in making many products for the machine vision market profitable. However, recent years has seen a few notable exceptions.

The lighting and optics has long been widely recognized as an essential first step to a successful vision application.[7,8] A clean image can make the difference between an easy to analyze scene, and an unreliable failure. Parts are not designed like optical elements, with well defined optical characteristics. They are designed as bearings, pumps, fasteners and any number of other items, the appearance of which is at best secondary. However, the optical performance of the part has a major impact on the ultimate performance of the machine vision system

There are many lighting techniques that have been developed both by accident and by design that can be used in a vision system. The objective of the lighting is to make the features of interest stand out in the image. Typically, this means the features of interest must be evident in a black and white image. Black and white cameras still provide the best resolution, cost, and flexibility for machine vision. Even for colored parts, separating colors with specific color filters typically provide better control that what can be realized with a color video camera (which, after all, are limited to red, green, and blue colors of wide wavelength regions).

Determining what is required in the image outlines the task and defines the limitation of performance that can be expected from the viewing system. A simple shape identification task may not need the high resolution needed to accurately gage a small feature to small tolerances. As the resolution performance of the lens system is approached, the lens degrades the contrast of the image until the small dimensional changes are washed out of the image. An initially low contrast image produced by the lighting further degrades the limiting resolution of the viewing system.

There are many other considerations in a machine vision application such as mechanical vibrations, fixturing, and space limitations. However, getting the lighting and optics right goes a long way toward a successful application. To facilitate getting the right image for machine vision, there have been a wide range of tools that have appeared on the market in the past few years. On the lighting side, some of these tools have included:

- Diffuse Lighting modules, both on axis and surrounding like a tent that help to decrease the sensitivity to local shiny spots or surface irregularities
- Directional Lighting modules, including line illuminators, dark field illuminators, and collimated illuminators that highlight surface textures, some point defects such as scratches, and surface irregularities such as flatness
- High Frequencies lights to permit asynchronous image capturing
- Highly Stabilized lights and high light uniformity, both of which provide for a more repeatable image

But the most visible explosion of systems evident in the trade show has been the use of LEDs, both with colors and white, to make new versions of almost all of the modules and layouts used by machine vision in the past, as well as some new designs tough to make with other lights (see Figure 2). The improved brightness and longer operation life of LEDs in the past years have made them a good alternative to halogen lamps that put out heat and have limited lifetimes or fluorescent lamps that remained bulky and expensive to customize. Although not yet powerful enough to be used in many larger field-of-view applications, and still a more expensive option incandescent lights, LEDs offer a degree of flexibility of design, and a potential ruggedness of manufacture that has otherwise eluded the machine vision designer in the past. Making an odd shaped line, a partial dome, or a surrounding light that can be adjusted by quadrants are all tools now in use in machine vision using LEDs as the light source. To illuminate a large part with a thousand LEDs may not be desirable. However, as the great majority of machine vision applications today are related to electronics such as IC chips, connectors, and other small components, the size issue had not been a big limitation. Quite a few companies have found the market viable for the manufacture and sale of LED lights. As LEDs make their way into more consumer products such as cars and appliances, this market is likely to both consolidate to stronger companies, and likewise grow and become more cost competitive.

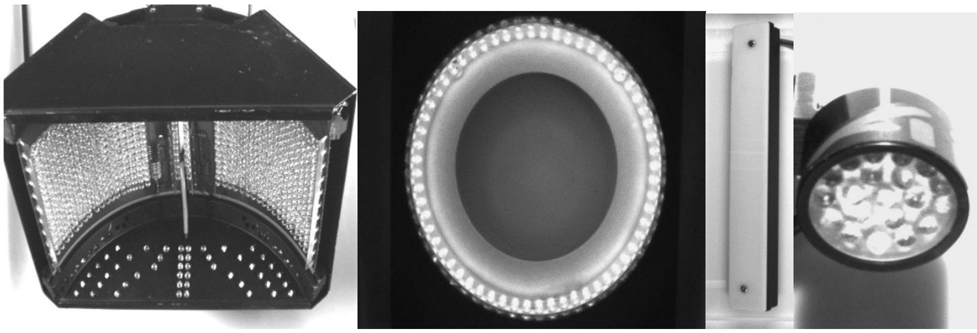

Light Tent Ring Light Line Spot

Figure 2. Some new lights showing flexibility of design using LEDs

In the optics area, manufacturers have also found it worth their while to address the machine vision needs, primarily on two fronts. The first area is just the availability of new high resolution lenses with better corrected fields. In the earlier days of machine vision, it was often argued that the ìsoftwareî could ìcorrectî the image. However, these corrections are something that takes up processing time, computing power, and in many cases are only partially successful. The technology for better image correction has been in place for a long time in such markets as 35 millimeter camera lenses, enlarger lenses, and micrographic optics. These non-video lenses have seen wide use, especially in small part (electronics and fittings) applications, taking advantage of special adapters to permit these lenses to be mounted on c-mounts of video cameras. Now, several companies have introduced lines of small, highly corrected, c-mount lenses that reduce aberrations such as field curvature and distortion, errors often up to 10% of the field in security camera lenses, by factors of ten or more (such as the lens data shown in Figure 3).

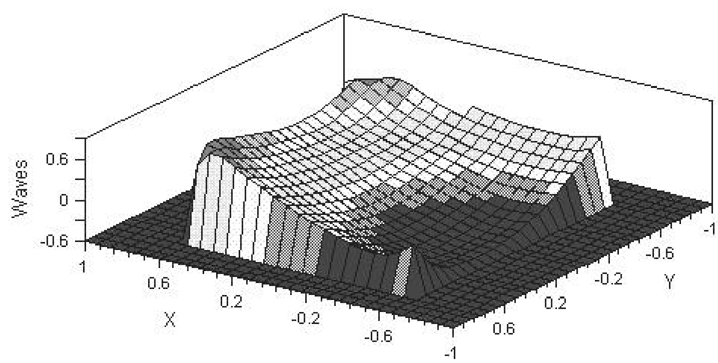

Figure 3. The corrected field of a high resolution c-mount lens, with less than 0.1 percent distortion.

The other advance in the imaging area has been the availability of a new range of telecentric optics. Telecentric lenses provide a means first to present a uniform intensity image onto the camera if the subject is uniformly lit. This may not sound like much, but in fact most common lenses vignette the light that is not in the middle of the image field, causing the light at the edges to be reduced. Just looking at an image, most people might not even see this effect as people are very tolerant to light level changes. However, given a vision system that is thresholding the light at some level in order to do the inspection, this light level variation is at least a nuisance. Even with adaptive thresholding, the limited dynamic range of cameras and the processing time taken for such filtering are both things that would be nice to not use up on something like light uniformity, given that the light sources themselves have become much better at presenting a uniformly lit field. The second benefit telecentric optics provide (the ones telecentric in object space) is produce a image that does not change in magnification for small shifts in the object distance (shown in Figure 4). This means the system can be much more tolerant to shifts in the position of the part under inspection, both within the field and in distance, without the character of the image changing, either in terms of perspective view or magnification. Telecentric optics do have their limitations. To be completely effective, there is some loss of light due to larger f-numbers, and the lens system does need to be larger than the field under inspection. But once again, with the great majority of the applications being in the electronics and similar small parts inspection area, telecentric optics have become a commonly used tool today.

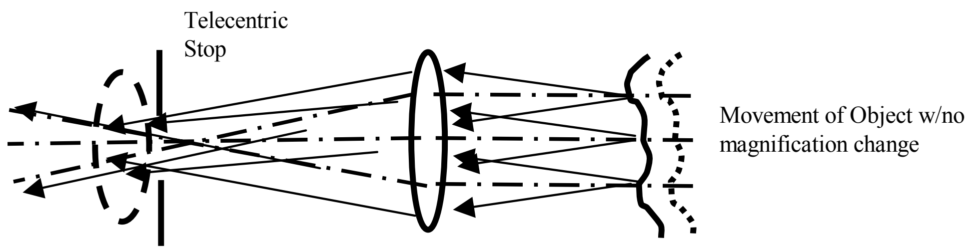

Figure 4. A telecenctric optical system collects a constant perspective, magnification and cone of light across the part.

2.2 Cameras

The next piece in a machine vision system is the camera. In the early days of machine vision, the application engineer was stuck with cameras made for closed circuit video, or needed to pay very high prices for higher quality cameras. The typical analog camera provided between 50 and 80 usable grey levels, with a varying degree of noise that was often different from one camera to the next. But what was probably worse from the perspective of machine vision, is the slight errors in what was actually a pixel or picture element. An analog camera had pixels on the video sensor chip, usually made to a high linear accuracy. But then the read out signal is turned into an analog voltage, that is digitized. The actual digitized pixels might be slightly different than the physical ones, a phenomena referred to as pixel jitter. This means that even if the vision system produced a mapping of the variations in pixel responses, the jitter could still cause an error in the voltage interpreted from a pixel due to this jitter. Synchronizing to pixel clocks and like techniques helped a bit, but still at the cost of more signals to deal with, and more computer overhead. The other factor with the old CCTV cameras was that they were limited to around 750 by 488 lines of resolution. This often meant more cameras to cover a given size field, just to obtain the resolution needed to do an inspection. The cameras were not expensive, but this added a need for multiple inputs on vision systems, more processing and the potential for errors from one camera to the next that were not corrected.

Today, the consumer market has helped to push the digital revolution to the old camera industries. Many vision systems today use, or at least have the capability to use, digital cameras that have provided better stability, higher dynamic range, and more pixels. Cameras with over 1000 by 1000 pixels of resolution are common place today. A camera offering 10 or 12 bits of pixel depth, that is, a light range of over 1000 to one (even allowing for a few counts of noise) are within the price range of many machine vision applications. This has made it easier, and more reliable to cover larger fields or obtain better resolution without complicated multi-camera systems. The advent of new communication options for cameras, including firewire, high speed internet, USB and the camera oriented camera link (which offers very high image speeds) has also made it easier to install cameras, network them together, and ultimately collect more data. The camera link standard used by some vision companies today will permit image speeds of over a hundred images a second, providing additional speed to inspection systems that otherwise would need multiple cameras running at 30 frames per second.[9] It is likely that over the next few years, the consumer market will drive the price of even the higher end cameras, up in the few thousand dollar range today, down to levels of the former analog cameras.

But even with all these great advances in the camera options of digital interfaces, more pixels, and lower noise, the big impact of growing digital capabilities today is in the form of ìsmartî cameras. A smart camera combines together a video camera with an on board processor and memory to create a vision system that can fit in the palm of your hand. Smart cameras in one form or another have been around for some time, going back to early systems that were very simple in rather clunky boxes. These early systems could typically do only one of a handful of operations at one time. Simple edge detectors that could find the distance between two edges were useful in alignments or gaps in assembled parts. Basic blob recognition provided simple optical character verification, or for that matter any simple pattern verification. But these systems were so limited, they were often little more than high end bar code scanners, but at a much higher price.

Today, smart cameras have a much larger range or operations, using more memory and processing power (typically pentium class) than older desktop computers had even 10 years ago. The types of operations available on a typical smart camera system can include[10]:

- Part location in position and rotation (allows for part movement or position variation)
- Analysis of multiple edge locations including counts, separations, and angles between edges
- Blob analysis to match complicated patterns, including doing full optical character reading (not just verification)
- Providing a wide range of outputs ranging from simple logic outputs to detailed numerical reports of fits to tolerances, amount of errors, and statistical information

In many cases, these cameras are made to go onto internet connections to allow them to be networked together. Since most of the processing is local, only the results or daily reports need to go over the network, removing the need for separate dedicated computers.

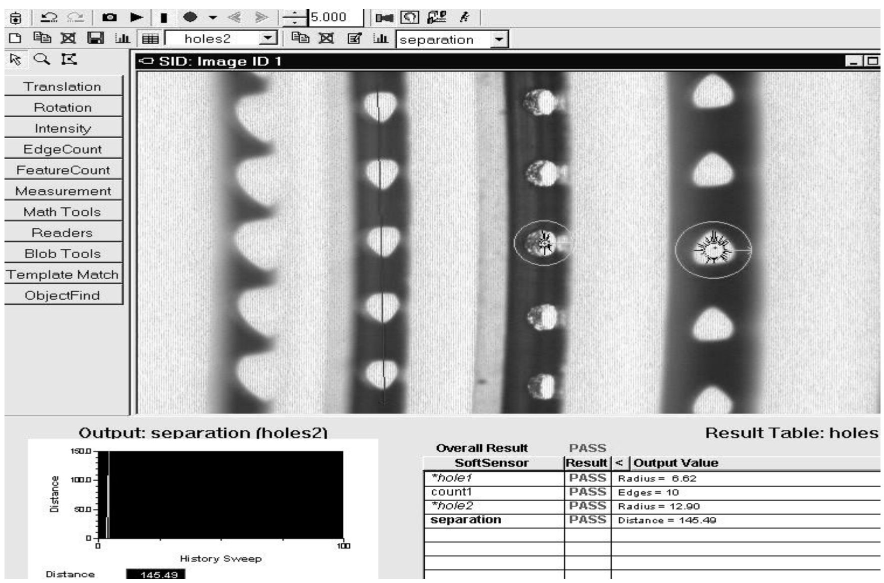

Figure 5. Typical smart camera applications of edge measurements and feature counting.

The degree of sophistication and cost of these smart cameras is still fairly wide, ranging from modern versions of the simple pattern matchers, now in the thousand dollar range rather then ten thousand, to full systems costing a few thousand. In general, the software is made user friendly, using pull down menus and icons to set up applications rather than C-code and low level communication protocols (see for example, Figure 5). However, these systems do not do it all. More complicated operations like image preprocessing, morphological operations, Fourier analysis or similar mathematical analysis and correlations are typically beyond what is reasonable to do with these smart cameras. However, with the capability today, there are a vast number of i goodî vision applications, things vision can do well, such as basic part measurements, simple inspections, hole finding or counting and the like that are very reasonable to do, both in effort and cost using smart cameras. What this means to the industry is many of those applications that just could not be justified because the vision system itself was going to cost thirty to sixty thousand dollars, can now be done with a smart camera at an equipment cost of a tenth of these amounts. It is still important to do a good application engineering job. Contrary to the claims of some salesmen, as with any system going into manufacturing, it is going to work much better, and be more reliable if engineered properly by someone trained to do so. But these applications do not demand the teams of software programmers, electrical engineers, and other developers needed in the past for just about any application of machine vision.

2.3 Software

This brings us to the third basic building block of machine vision, the software. Not too many years ago, to set up any vision application meant either being a good programmer using C or similar language, or worse yet, learning a dedicated vision language, built around mnemonics, command strings and controller codes. Basic C-code libraries are still available today, and in fact have become fairly inclusive providing a wide capability for the programmer. However, for most applications, vision systems offer a user interface with access to all of the functions, filters, communications and the like to the operator without writing a single line of code. In many software packages, the user interface is made as simple as possible, emphasizing application oriented operations, with only a limited amount of ìfiltersî that are not always fully defined for the operator to use. Typical application operations represented by simple icons might include:

- Read of bar or 2D matrix code
- Measure the separation or angle between two edges
- Measure the perimeter, volume, diameter of some feature
- Measure the separation of two features (such as hole according to their centers)
- Find and locate a particular pattern (often registering other measurements to datum locations)
- Find any changes in the scene patterns (such as defects of a certain size)

The icons provided often suggest common tools a shop person would recognize, a caliper, a protractor, or a tape measure. In these systems, any filtering is general preset by the software, or has simple settings (like 0, 1, or 2) to employ some appropriate filter for that type of operation. The result is often displayed as a block diagram or flow chart that can be easily followed by the operator. In some systems, the user interface will generate an executable file (often in C or C++) for use in operation, just to provide better speed of operation. Simple additions of setups desired are often available through small scripts, usually based upon visual basic, with clear examples and instructions for the end user.

Many fewer vision systems use any form or proprietary hardware anymore. They may use multiple processors, high speed memory and some fast graphics, but it is not often that the end user needs to know how to do low level programming of any type to get the system to work properly. Even the camera setups, which at one time often required the writing of special configuration files, is often plug and play using the standards mentioned above. The high speed of todayís computers has replaced the high end workstation costing a hundred thousand dollars, with a simple system selling for under five hundred. One variation that has seen some popularity in the vision industry in the PC104 computer platform. This platform allows a full operating system type computer, with hard drive and (external) CD drives to fit into a small case the size of a portable CD player. Such a system can offer more performance than a smart camera, particularly in terms of longer term memory, in a package that can be built into the vision box or mounted under a conveyor belt without much problem.

The user friendly nature of machine vision software has brought the setup, programming or changes into the capabilities of plant engineers without extensive programming experience, as well as many shop floor maintenance personnel who can now much more easily maintain systems that at one time could only be serviced by a high level programmer not typically available in a plant. Parts programming rarely needs to be done by the vendor anymore (though most are willing to do so). This puts the control and schedule of new part introduction, tolerance changes, or the ability to add new checks to diagnose manufacturing issue within the hands of the production manager. This ease of use, along with the lower prices, has made vision systems more attractive to many industrial inspection, alignment, and simple gaging applications that in the past may have been too expensive both to purchase, and to maintain and setup for new parts.

3. 3D MACHINE VISION: THE EMERGING FRONTIER

So far, the discussion has been very general but primarily focused around 2D machine vision. Two dimensional machine vision relies upon what can be seen or inferred from a 2-dimensional image, be it flat dimensions, patterns, or defects, for use in controlling or monitoring the manufacturing process. However, many specifications, defects, and parts in general are three dimensional (3D) in nature. The digitization of manufacturing has not only streamlined design of new parts, in many cases it has allowed designers to think in three dimensions. Multi-axis machine tools, EDM or ECM (electrical discharge or electro-chemical machining) have the capability to make complex 3D surfaces. CAD (computer aided design) systems also have 3D capabilities today, able to produce surface shapes and solid models not available to the designer of the past. This affords new opportunities for ì virtualî assemblies, aerodynamic shapes, and more effective surface designs that were very hard to even visualize in the past. These new designs, as well as 3D shapes and relationships in many assemblies today need faster, and more accurate tools to inspect, verify and measure three dimensional shapes and relationships.

Just as with 2D machine vision, 3D machine vision offers the potential for that higher speed and greater flexibility demanded today.[11,12] Three dimensional feature verification is currently dominated by CMMs measuring a limited sampling of points, or for speed fixed gages made as dedicated tooling for each part. Together, the 3D part inspection base represents several billions of dollars a year in manufacturing expense. Just as with 2D machine vision, some measurements may not be practical with mechanical systems at all. The same market that dominates much of the rest of machine vision, electronics and small components and assemblies, is already dominated by 3D machine vision tools. Now, the technology of 3D machine vision is poised for the leap into wide spread production applications, just as 2D vision did 20 years ago.

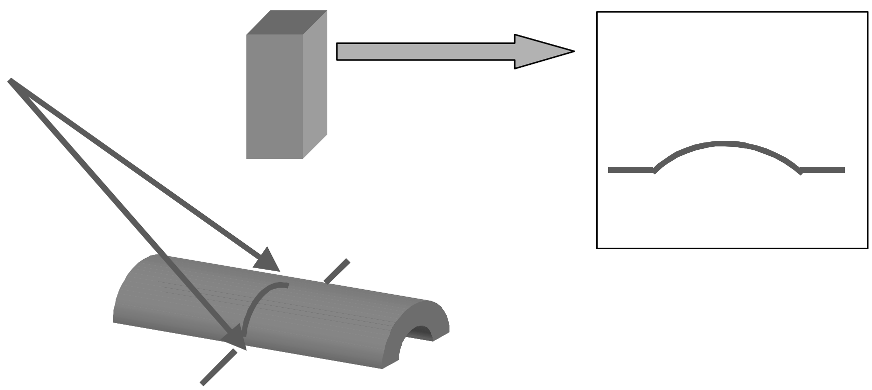

Figure 6. Basic laser triangulation system using lines of light.

So where is 3D machine vision today? The majority of systems in use today in the area of 3D machine vision fall into two application categories. The first is the dedicated sensor market, characterized by the electronics industry and similar specialty areas such as welding. These systems primarily consist of structured light based sensors that provide lines of 3D data such as shown in Figure 6.[13-16] Applications such as solder paste volume, verifying chip lead coplanarity, and general component placement have all been highly facilitated, often dependent upon specialized 3D machine vision systems, typically using some form of structure light to encode 3D, or at lease general shape information into 2D images that can be analyzed using the basic principles of triangulation. On small parts such as IC chips, these systems have been able to provide micron level 3D information at speeds approaching video analysis (hundreds of thousands of points per second, but typically limiting themselves to small areas, shallow depth ranges, or specialized applications (like chip lead coplanarity).

The other application category of 3D vision systems has been as a more *general purpose* tool, that in most cases is really just the sensor end of the application, generating at least multiple lines of data out to full 3D clouds.[17-21] That is, the system spits out 3D data points at rates of a few tens to hundreds of thousands of points per second, then leaves it as an exercise for the user to figure out what to do with then information. The two small in-roads to using full 3D data has been doing *simple comparisons*, either to models or other parts, or doing a form of *reverse engineering*, generating polygonalized or otherwise fitted surfaces than can be analyzed (manually) to extract CAD type of data to rebuild the design.

The most common general purpose application area today, reverse engineering, has been progressing regularly over the past 10 years with the help of a number of high end software packages made for that purpose. But as an operation generally limited to new designs from forms, tool verifications, or old part repair the application base for reverse engineering systems is limited. When digitizing a model or repairing a part, the operation is not typically a part of any production system, time is not of the essence and as such, 3D vision technology competes with automated CMMs, measurement arms, and today a wide selection of hybrid machines that offer some optical scanning, such as a line or point laser probe, along with the familiar convenience of traditional touch or mechanical scanning probes on a general purpose CMM platform. Almost every CMM manufacturer offers some form of optical probe option, ranging from laser probes to video cameras for speeding up operations like hole locations or area scanning to varying degrees of success. In addition, a number of dedicated manufacturers specialize in optical based CMM systems for surface scanning or video inspection equipment working as general 2D vision tool with the ability to also move in three dimensions and measure some depth information.

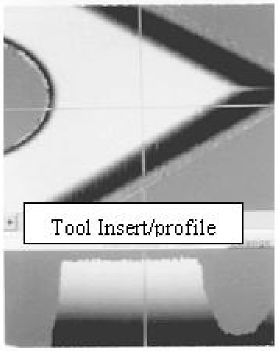

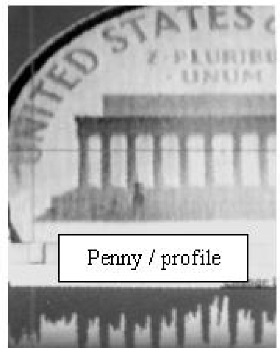

Figure 7. Examples of 3D gray scale depth maps, where the Z depth dimension is represented by grey level.

The other application of comparisons of parts to models is more in keeping with the traditional role of machine vision to monitor and quantify the manufacturing process. Unfortunately today, the measurement with a 3D vision sensor for part analysis typically leads to color encoded pictures. However, there are no part specifications that related to color encoded pictures, which in general are just qualitative visualizations at best. What is needed for inspection is typically numbers, of specific features, defined in specific ways, and outputted in a format easy for interpretation by the operator. This later advance, the quantification and interpretation of output, is just beginning in the 3D industry today.[22-25]

Creating the tools of data output was a big stepping stone that occurred in 2D machine vision almost 20 years ago. The challenge posed by 3D vision systems is the volume of data vendors are trying to present, not really the interpretation. Just as a 2D vision system digitizes every pixel in the image, so do 3D vision systems. What the 2D vision industry learned is you do not, and do not want to use every pixel of data. Even parts that have 3D shapes defined, they are currently most typically defined in terms of cross sectional fits that is not unlike any 2D outline or fit. Most other features on a part are defined by simple relationships such as angles and distances, just like in 2D vision, but now with another degree of freedom. What the manufacturer needs is numbers quantifying relationships, distances, sizes or locations of features, ideally as compared to nominal values and tolerances. Any 2D vision system will provide this type of capability. In many cases, the same analysis used in 2D vision could be used on 3D data by treating the grey scale depth map generated by the 3D system as if it were a 2D image.

So why not just let the vectors of distances and shapes be defined in 3D rather than a single plane? A measure of distance, angle, or diameter is still defined the same. Almost without exception, almost no 3D vision system has this same capability. The best that is typically offered is to export the full, quarter million to 10 million set of X, Y, Z points to a high end 3D analysis package which can cost as much or more than the 3D vision system, and require extensive training to use. The technical issue seems to be in defining 3D datums, orienting the tools of 2D analysis by registering their positions and orientations in 3D space, rather than only in 2D space. At one time not too many years ago, this was a problem because of the processing time to execute 3D matrix transformations similar to those used in 2D systems. But just as with 2D system, locating and registering an image used to be a big thing, but with current computing power, is routine, available in the simplest of smart cameras. This can also be done in 3D, it just hasnít happened as yet. The change that drove this in 2D vision technology was investment, volume applications that could justify development of tools, and growth of vision companies to have sufficient resources to execute the needed development. There are billions of dollars spent each year on fixture gages and manual measurements of 3D features on manufactured parts. Hopefully, at some point, the investment will be made to allow this market to be addressed

To address this potential bigger application base of 3D vision, namely production monitoring and control, there are already a wide range of systems with the general capability to obtain large amounts of 3D information, at the speeds needed (typically seconds to something under a minute), and at resolutions that in just the recent couple of years starts to fit the thousandth of an inch (1 mil) level of inspection and measurement capability demanded by machining and forming manufacturing operations today. With some volume, just about any of these systems would reduce in price to attractive levels. Even at current typical prices which range from the forty thousand dollar level for the 3D equivalent of the basic smart camera (basic resolution of a few mils, 640 by 480, few seconds to collect data) to the higher end, finer measurement systems costing one hundred to two hundred thousand dollars (typically better than one mil, over 1200 by 1000 pixels, and few seconds to under a second) which would compare with a high end 2D vision processor selling for thirty thousand, there are still many applications that could show a payback of 6 months, if the same analysis tools were available for 3D vision as 2D vision already has available.

Table 2. Typical capabilities of current 3D vision technologies today.

SENSOR TYPE	RANGE RESOLUTION (inches)	SPEED (Points/second)	SURFACE TYPE	POINT DENSITY (pts/inch^2)
Point Scanner/Triangulation	0.0001	1500	diffuse	1,000,000
Point Scanner/Conoscopy	0.00005	800	diffuse to shiny	25,000
Point Scanner/Interferometric	0.00001	200	smooth	1,000,000
Line Scanner/Triangulation	0.003	150,000	diffuse	25,000
Line Scanner/Conscopy	0.0005	4000	diffuse	10,000
Line Scanner/L. Radar	0.04	50	diffuse	500
Area/multi-line	0.003	150,000	diffuse to shiny	3,000
Area/Phase shifting	0.001	120,000	diffuse to shiny	4,000,000
Area/moire	0.0001	300,000	diffuse to shiny	1,000,000
Area/radar	0.04	10,000	diffuse	500
Area/stereo	0.005	5,000	textured	500
Area/Interferometric	0.00001	120,000	smooth	1,000,000

4. APPLICATIONS

As with any tool, there are good applications of machine vision where it affords good capabilities, and there are poor applications that can best be done by other methods. The specific attributes machine vision has that best distinguishes the application domain include:

- non-contact, so low change of damaging delicate parts
- fast data collection, where many point may be needed say for SPC applications
- line of sight part access is needed, if you can't see it, you can't measure it with vision
- interaction is optical, meaning it is more important how a part looks, that how it feels

These attributes are neither inherently better nor worse than the attributes of other types of gages, but they are different. Machine vision has been used in very hostile environments, such as steel foundries or nuclear reactors. But machine vision is not used the same as a mechanical gage. Stray light, reflections and glints will affect machine vision but not a mechanical gage. Extreme temperatures and mechanical shocks will affect mechanical gages, but not machine vision. This means that the way a vision system is applied is different that how a different technology is used. One of the biggest mistakes in applying machine vision as a tool is trying to apply it as though it was some other technology entirely. A screwdriver makes a poor hammer and a worse saw, yet these are perhaps more alike than machine vision and mechanical point probes.

There are some generalizations we can make about what makes for a good potential application of machine vision today, and what makes a bad one. There will always be exceptions to these generalizations, but it provide some guidelines. Some typical good parameter of a machine vision application may include:

- a large standoff to the part is needed, an assembly with an active robot in the way
- the environment is mechanically hostile, such as in a hot steel rolling mill
- part touching is a problem, such as a fragile ceramic part
- many small feature must be measured, such as circuit boards
- the part interacts with light predictably, such as fluorescing grease

By the same discussion, there are a number of features of an application which may make machine vision a poor choice given the capabilities of machine vision today. Some typical bad application parameters for a machine vision application may include:

- the part appearance varies widely, but we don't care about that
- there is poor access to the features of interest, so they can not be seen
- the air is difficult to see though
- only a few points are needed, and the current mechanical/electrical sensors work fine

Machine vision need not be a one to one replacement of something that is already being done by other means. There are plenty of applications which were not done at all in the past, because other methods where not able to due them. Today, they are seen as good applications of machine vision. Some examples include pill inspection for defects, electronics component and connector inspection for bent or missing leads, and shape verification of gaskets. With technologies like smart cameras, many more of these applications now make economic sense to attempt with machine vision technology. Machine vision today plays many roles in manufacturing, such as go/no-go inspection of small plastic parts using in circuit breakers, visual defect detection on paint finishes of consumer product like toys, assembly verification of motors in washing machines or wires in cars, robot guidance for welding or assembly and gaging of bearings. Although they are closely related, each of these tasks are fundamentally different. Each places different requirements on the system's processor, optics, mechanics and, in some cases, even the cabling.

A machine vision system today can accommodate a degree of misposition, both in rotation and translation, and still make a reliable measure relative to the part. The critical factor in fixed mechanical gages is the fixturing. The critical factor in machine vision is contrast, presenting an image which contains the dimensions of interest, regardless of where it may be. Highly uniform light sources such as LED arrays and telecentric optics has eased this task for machine vision developers. This is not to say fixturing is not important to machine vision, processing time can be spent locating a part. Ten of fifteen years ago, this part position finding could take several seconds. Today, a small misposition rarely takes more than a tenth of a second to correct. This degree of flexibility if one of the key features of machine vision as a gaging tool.

In general, there are many areas of precision gaging that call for metrology lab precision, that is just not viable with most 2D or 3D machine vision today. However, there are many on-line applications where machine vision can provide valuable information for control of the process. One of the more popular instances of the use of machine vision has been in the fit and finish of auto bodies. A consortium effort between industry and university workers used machine vision gaging to control the fit of auto assembly to better than 1 millimeters using vision technology. This was a substantial step toward a consistent level of quality for the manufacturer. Another important area has been the inspection of web products such as paper, plastics, and primary metals for small surface imperfections. Smart cameras are being used for many applications formerly done by manual optical comparators, such as hole sizing, and basic outline tolerancing. Three-dimensional vision technology is routinely used for electronic component placement guidance. Fast gaging of forgings and stampings is just starting to make headway against hard gages and CMMs, but with fast computing power is more attractive than ever today as a option. A 3D system capable of measuring a foot size volume at thousandth of an inch accuracies in a few seconds is in the same price range as a good automated CMM ($100k+ range), but for complicated parts, can be over 1000 times faster than even a scanning CMM. A summary of what can typically be expected from machine vision systems, from smart cameras on up, is summarized in Table 3.

Table 3. Typical gaging capabilities for machine vision in 2D.

OPERATION	PART SPEED	# OF FEATURES	RESOLUTION
edge location	30/second	30 per part	1/2000 FOV
hole center/size	30/second	5 per part	1/10,000 size
feature location	15/second	10 per part	1/1000 FOV
3D flatness	2 to 10 seconds	1000 points	0.0001 inch
3D shape compare	5 to 20 seconds	40K points	0.001 inch

Economic considerations must always be addressed when designing any production systems. Economic considerations for machine vision applications include:

- How important is the inspection function?
- What is the relationship between the cost of capital, the cost of servicing it, and the cost of labor?
- What are the long-run demand trends?
- What are the costs in a competitive market of not being automated?
- What would be the cost of any retraining?

For example, perhaps you make bearings to go into engines. The inspection you might do is look for an oil groove. If you do not inspect for the oil groove and it is missing, the engine may fail. The initial capital for automating the inspection and the retraining cost may be as much as a year of labor for manual inspection. If this part may soon be obsolete, then the automation may not pay for itself. However, if the part is around for the long haul, there is a payback time which some may consider long. However, we must also consider the cost of charges back for failed engines, and the competition. If the competition can produce a better part without this flaw (through whatever means), then you stand to loose 100 percent of your business. So in todayís market you must also ask how much is "the business" worth when deciding on the cost of installing a new tool like machine vision. This is the real potential payoff of machine vision tools today.

5. THE FUTURE

In the past 20 years, machine vision has grown from an ugly step child status filled with wild claims, to an accepted tool in many areas of manufacturing. The current market for machine vision is several billion dollars a year today (according to surveys of the Automated Imaging Association), with the potential to more than double in the next five years. The acceptance of machine vision, as with any new technologies, has been slow. Machine vision is different from what manufacturers have used in the past, but it is a tool very in tune with the modern thinking of computer integration in the factory. There is a change occurring in how process control is approached, and how the production person thinks about data collection. Industry is no longer satisfied with throwing out bad parts, it doesn't want to make bad parts at all! Most manufacturers have typically not known much at all about such optical technologies, but they are learning.

Just a few years ago, machine vision was viewed with suspicion as some form of black magic, tied into equally magic computers. Just as computers have become accepted in the work place, so has machine vision. Even 20 years after the wide spread first uses of machine vision, there are still new applications coming about for machine vision each day. Although machine vision can do much more than it could twenty years ago, it does not do everything. There are some, off-the-shelf applications of 2D vision technology, in standardized systems doing high volume inspections from bottles to sheet product. Machine vision ìstandard productsî that inspect paper, printed labels, pills, painted surfaces, and plastics can be purchased practically off the shelf today. Although there are still plenty of new applications being tried every day, not every application is necessarily a new development effort as it was in the past. Much of the development on these early applications has been done, documented, and proven out in practice. Even with this beginning, machine vision systems have not reach the commodity status. However, with the foothold of general acceptance growing for machine vision there is now a new realm of consumer related products envisioned, ranging from security systems to vision in cars to read street signs to home units to watch the baby.

Over the past several years, there has been quite a bit of consolidation in the vision industry.[26] Companies have combined with others to broaden their market, adding defect inspection to part sorting or web inspection to glass container lines. They have also consolidated to access new markets, combining developments for the US, European, and Asian market into global companies. These larger companies have better resources, can do more research to improve their product, and can better market to a global market. With computing advances and consolidation providing cost savings to reduce prices even more on low cost systems, the new tools such as smart cameras and LED lights will be able to into areas of inspection formerly done by hand or machine operators, not as a person replacement, but as a quantitative tool to verify subtle features not discernable by the human eye. The already low prices of some of smart camera systems will continue to drop, making ëautomaticí part tracking using OCR, 2D bar codes, and even part features a much more robust replacement for routing sheets and hand held bar code readers.

Specialized areas of 3D machine vision such as three-dimensional contouring are just now starting into industries with great potential for the future. Clearly the jump needs to be made from the 3D sensor creating pretty color images, to the measurement and inspection tool providing tolerance compliance information. The move from sensor manufacturers of 3D sensor, to producers of 3D based inspection and gaging systems will require investment, and likely consolidation within the 3D vendor market, just as has been occurring in the 2D vision industry. It is quite possible that 3D copiers will be available in less than 10 years. But it is not likely to be the small 3 to 5 person 3D sensor shops or today providing those 3D copiers. To produce a commodity 3D copier will require design engineers, CAD engineers, user interface designers, production engineers and a wide range of other talents that may have little to do with that pet 3D technology embraced by the founders. 3D copiers as a widely used "commodity" item could drastically change the way we think of manufacturing, driving manufacturing toward being a truly flexible operation. The momentum is there, and the potential is quite interesting.

6. REFERENCES

1. *Science, Technology, and the Issues of the Eighties: Policy Outlook* for the American Association for the Advancement of Science, A. Teich and R. Thornton Editors, p. 7, Westview Press, Boulder, Colorado (1982).

2. . K. Harding, ìThe Promise of Machine Vision,î **Optics and Photonics News**, p. 30 May 1996 Optical Society of America

3. . *The Machine Vision Market: 2002 Results and Forecasts to 2007*, Automated Imaging Association, Ann Arbor, MI, (2003).

4. *McGraw Hill Dictionary of Scientific and Technical Terms*, 3rd edition, McGraw Hill, New York (1984).

5. Nello Zuech., *Applying Machine Vision*, , John Wiley and Sons: New York (1988).

6. *Automated Visual Inspection*, B. Batchelor, D. Hill, D. Hodgson Editors, p. 10-16, IFS (Publications) LTD, UK (1985).

7. K. Harding, ìMachine Vision - Lightingî, **Encyclopedia of Optical Engineering**, Marcel Dekker, New York (2003).

8. Eric J. Sieczka and Kevin G. Harding, "Light Source Design for Machine Vision", **SPIE Proceedings** Vol. 1614, **Optics, Illumination and Image Sensing for Machine Vision VI**, Boston, Nov. 1991.

9. John Vieth, ìCamera Interfaces Vie for Vision Applications,î **Vision Systems Design**, Sept. (2003), p. d3.

10. K. Harding, "Machine Vision Based Gaging of Manufactured Parts," **Proc. SME Gaging '95**, Detroit, Nov. 1995.

11. L. Bieman, "Three-Dimensional Machine Vision," **Photonics Spectra,** May 1988, p. 81.

12. E. L. Hall and C. A. McPherson, "Three Dimensional Perception for Robot Vision," **SPIE Proceedings** Vol. 442,

p. 117 (1983).

13. M. R. Ward, D. P. Rheaume, S. W. Holland, "Production Plant CONSIGHT Installations," **SPIE Proceedings** Vol. 360, p. 297 (1982).

14. G. J. Agin and P. T. Highnam, "Movable Light-Stripe Sensor for Obtaining Three-Dimensional Coordinate Measurements," **SPIE Proceedings** Vol. 360, p. 326 (1983).

15. G. L. Oomen and W. J. P. A. Verbeck, "A Real-Time Optical Profile Sensor for Robot Arc Welding," **SPIE Proceedings** Vol. 449, p. 62 (1983).

16. Kevin Harding and Dale Markham, "Improved Optical Design for Light Stripe Gages," **SME Sensor '86**, Detroit.

17. K. Harding, ìOverview Of Non-Contact 3D Sensing Methodsî, **Encyclopedia of Optical Engineering**, Marcel Dekker, New York (2003).

18. M. C. Chiang, J. B. K. Tio, and E. L. Hall, "Robot Vision Using a Projection Method," **SPIE Proceedings** Vol. 449, p. 74 (1983).

19. J. Y. S. Luh and J. A. Klaasen, "A Real-Time 3-D Multi-Camera Vision System," **SPIE Proceedings** Vol. 449, p. 400 (1983).

20. K. Harding, "Moire Techniques Applied to Automated Inspection of Machined Parts," **SME Vision '86,** Detroit June 1986.

21. K. Harding and L. Bieman, "Moire Interferometry Gives Machine Vision a Third Dimensional," **Sensors**, October 1989 p. 24.

22. N. Kerkeni, M. Leroi, and M. Bourton, "Image Analysis and Three-Dimensional Object Recognition," **SPIE Proceedings** Vol. 449, p. 426 (1983).

23. J. A. Beraldin, F. Blais, M. Rioux, J. Domey, "Signal Processing Requirements for a Video Rate Laser Range Finder Based Upon the Synchronized Scanner Approach," **SPIE Proceedings** Vol. 850, p. 189, (1987).

24. H. Ailisto, T. Heikkinen, A. Kilpela, R. Myllyla, M. Pietikainen, "Acquisition and Processing of Range Data Using a Laser-Based 3-D Vision System," **SPIE Proceedings** Vol. 850, p. 174 (1987).

25. D. J. Svetkoff, P. F. Leonard, and R. E. Sampson, "Techniques for Real-Time Analysis of Laser Radar Data," **SPIE Proceedings** Vol. 521, p.302 (1984).

26. Nello Zuech in ìEye on Imaging ñ Views and Opinions: State of the Machine Vision Marketî **Coreco Imaging Newsletter**, www.imaging.com, Sept. 24 (2003).