

TAKING CONTROL OF MACHINE VISION

IMPROVEMENTS IN PROCESSING POWER AND SOFTWARE AVAILABILITY ALONG WITH THE RISE IN DIGITAL CAMERAS ARE MAKING THE USE OF MACHINE-VISION SYSTEMS FOR INDUSTRIAL CONTROL AN INCREASINGLY ATTRACTIVE APPROACH. THESE TRENDS ARE SIMPLIFYING SYSTEM CREATION, BUT GETTING THE SYSTEM TO FUNCTION AS INTENDED STILL REQUIRES CAREFUL ATTENTION TO DETAILS.

BY RICHARD A. QUINNELL • CONTRIBUTING TECHNICAL EDITOR

In humans, sight is one of the most important senses. Vision allows us to identify objects, examine them without touching them, determine spatial dimensions and relationships, and navigate safely through our world. The same benefits accrue to the machines equipped with vision systems, when the camera and processing are configured to match the application's needs.

Numerous applications currently employ machine-vision systems, many of them in factory automation. Vision-equipped machines inspect products for flaws, identify and sort objects, measure dimensions, and align and position materials for automated assembly. Nonfactory applications include vision-based navigation

for autonomous vehicles and safety systems for automobiles. As diverse as these applications are, however, they share three common elements: a camera, a processor, and image-manipulation software.

Not surprisingly, the camera is a key element in establishing a machine-vision system's capabilities. The camera

sets the level of detail, or resolution, that the system can distinguish. It also sets an upper limit on the system's frame rate, or the speed at which the system can generate images, and the shutter speed or image-capture time. Frame rate dictates how quickly the control system can obtain updates, and shutter speed affects how quickly objects can be mov-

ing through the field of view. In manufacturing systems, these factors control the manufacturing throughput that the system can handle.

Both analog and digital cameras are available for machine-vision systems. Many high-end analog cameras, however, are designed to form television images and so offer few options for either frame rate or resolution. The advent of Internet video has made available some low-cost analog cameras with other resolutions, typically VGA or fractions thereof. Such cameras may be suitable for applications that are highly cost-sensitive, but they offer limited performance. All analog cameras require the use of a separate digitizer—usually part of the frame-grabber hardware—to capture the image for processing.

Digital cameras, by their very nature, do not need external digitizers and have the additional advantage of being free from the resolution and frame rate of

video standards. The explosion of digital cameras in consumer applications, from photography to cell phones, has stimulated significant advances in the technology during the last decade. Digital cameras are now available in a range of resolutions with an equally broad selection of achievable frame rates, at ever-decreasing cost. Most machine-vision-control systems use digital cameras as their "eyes."

SELECTING A CAMERA

There are a variety of trade-offs to consider in selecting a digital camera. One of the most significant is resolution ver-

AT A GLANCE

- ▶ Machine vision simplifies system control based on size, shape, color, and position of objects.
- ▶ Digital cameras offer trade-offs between resolution and frame rate.
- ▶ Processor options include CPUs, DSPs, and dedicated hardware at the component level, with many board-level systems available.
- ▶ A new generation of high-performance image processors is targeting automotive applications, such as hazard-detection systems and camera systems that monitor drivers.

sus frame rate. In general, the higher the camera's resolution, the lower the frame rate it can achieve. This trade-off stems from the way in which the image sensor of the digital camera operates.

The core of a digital camera is a CCD (charge-coupled device) that is the image sensor. As Figure 1 illustrates, the CCD's basic configuration is a rectangular array of light-sensitive cells (picture elements, or pixels) that connect to transport cells, forming a row-and-column array. Light falling onto the array generates charges in the pixel. Command signals transfer the charges from the pixels to the transport array, where they move buck-

LIGHTING FOR MACHINE VISION

The ability of a vision system to gather accurate image information depends entirely on the light that enters the camera's lens. Inadequate or inappropriate lighting can compromise the system's ability to distinguish the details it needs for successful inspection, recognition, or measurement of the object being viewed. Although image-processing software can compensate somewhat for such factors as underexposure and overexposure, glare, reflections, and shadows can wreak havoc with many vision applications. And, if color is an important factor, the right lighting is essential for the accurate detection of color differences.

Whether you are using monochrome or color images, one important lighting parameter to consider is the stability of the illumination levels. These levels affect the brightness and contrast of the image and contribute to high signal-to-noise levels in the image. Often, ensuring illumination stability requires control over the environment. For example, some fluorescent lights change intensity 5 to 10% when the ambient temperature shifts as

little as 10%. Design of a machine-vision system should seek to control illumination levels if possible and adapt to changes when control is not an option.

Artificial-illumination sources that depend on ac electricity will inevitably flicker. One way to avoid the problems that this flickering can cause is to ensure that the ac frequency is higher than the machine-vision system's frame rate, so that the exposure time of any given image integrates several cycles of light variation. Too low an ac-line frequency lowers the throughput of such systems as automated inspection stations on an assembly line.

Geometry is another important lighting factor to consider. Geometry describes the alignment between the light source, the object being viewed, and the camera. The optimum geometry depends on the object's physical characteristics, which affect the reflection, transmission, scattering, absorption, and emission of light. Different surface textures, for instance, require different lighting.

THE RIGHT GEOMETRY

Figure A, which shows the results of two light positions on the image of a glossy surface with grayscale markings, highlights the importance of lighting geometry.




Figure A. Incorrect lighting in a machine-vision system can eliminate the visibility of the features upon which the system bases control (courtesy Edmund Optics).

An illumination source high above the surface can generate glare that compromises the image, washing out differences in the markings. Placing the source at a low angle to the surface and adding a diffuser to reduce specular reflections help in the detection of small visual differences on such surfaces.

Deeply textured products, on the other hand, produce significant shadowing when the illumination source is low to the surface. If the application is examining the texture, this approach may be appropriate. If the purpose is to look for marking on the surface, however, the shadows become, in effect, visual noise for the detection algorithms.

Other surfaces may require different approaches, so knowing the object's surface qualities is essential to the creation of appropriate illumination in a machine-vision system. Fortunately, guidelines are available. The CIE (International Commission on Illumination) provides recommended lighting geometries for a variety of surfaces through its Web site, www.cie.co.at/cie.

et-brigade style through the array to the charge sensor that provides the digital readout. CCDs from VGA resolution to more than 10M pixels in square or rectangular arrays are available.

A basic digital camera has only one charge sensor, so reading an image out of the camera occurs one pixel at a time. This one-at-a-time action ties the camera's frame rate to its resolution. For a given CCD-process technology, there is an upper limit to the speed at which the transport array can move the pixel charges. So, the larger the array, the longer it takes to read out an entire frame.

This trade-off is not absolute. Camera CCDs with multiple charge sensors are available. They break the image into nonoverlapping blocks that you can read out in parallel, increasing the achievable frame rate. However, balancing the mul-

tiplied conversions from charge to digital value so that the independent blocks will produce matching images complicates the design effort.

It is also possible to increase the frame rate of a high-resolution camera by using only a subset of the total image. Cameras that offer this feature allow the control system to specify an image area of interest for readout rather than shift out the entire array. This approach results in a smaller image at a faster frame rate.

APPLICATION TRADE-OFFS

The optimum balance between frame rate and resolution for a machine-vision system depends strongly on the application. High resolution, for instance, is necessary when looking for small defects on large objects or making precision measurements of object dimen-

sions. High frame rates help increase system throughput in terms of objects examined per unit time or the time needed to scan a large object.

The balance of frame rate and resolution affects camera cost. Fast multiblock and area-of-interest cameras are more costly than lower speed cameras of the same resolution. Higher image resolutions mean more expensive sensors and usually need more expensive optics in the camera to achieve the proper depth of field and field of view. Along with cost considerations, you must also factor into the camera choice such requirements as illumination (see sidebar "Lighting for machine vision") and color (see sidebar "When color matters").

The choice of camera sets the upper limit on achievable performance in a machine-vision system, but it is not the

WHEN COLOR MATTERS

Color is a subjective value that depends on an object's illumination and viewing environment as well as its spectral properties. Cameras do not see color, however; they detect levels of integrated spectral information that becomes a specific color only when a human observer views it under some set of lighting conditions. To automate color-based inspection, then, a machine-vision system must model the behavior of human eyesight.

Color in machine-vision systems typically has one of two functions. One is to use color as a basis for identification or discrimination of objects, as in a color-sorting system. Color-sorting systems need only to be able to distinguish among various, often quite different, colors, so their lighting and image processing are not particularly critical. A crayon manufacturer, for instance, can use a simple color-sorting

system with standard room lighting to ensure that there is only one red crayon in each box.

The other function of color in machine vision is to monitor the color itself to ensure that it meets production specifications. If a crayon manufacturer wants to guarantee that the red crayon in each color-assorted box looks identical to that same crayon in every other box, it needs a visual-color-matching system. Only color matching can determine whether two red crayons look identical.

COLOR MATCHING

Visual-color-matching systems must judge color in the same way that humans do. One way to make this judgment is to use the "golden-reference" strategy by choosing a unit that is as close to ideal as possible. The inspection system then uses relative-color analysis to measure visual-color differences between each item it ex-

amines and the golden reference.

The precision of color-matching systems depends in part on the illumination source's spectral characteristics. The light source needs to include adequate emission at every wavelength within the visible range—380 to 720 nm—to produce results that match human perception.

The system should avoid the presence of bright spectral lines in the illumination spectrum to be able to detect metamerism: two objects that look the same under one light source but different under another light source. Metamerism occurs because human visual systems can distinguish only broad ranges, not detailed wavelength information. The additional illumination at a specific wavelength that spectral lines provide can fool the eye into seeing a color that is not inherent to the object. Few cameras detect detailed spectral informa-

tion, so metamerism affects most cameras in the same manner as it affects humans.

When a system uses color matching solely for pass-fail inspections, the way in which the system quantifies the color is irrelevant; any color scheme will work. If the system is to quantify mismatches for the purpose of correcting the production process, however, the choice of color-description schemes becomes important. Standard color spaces such as the RGB (red-green-blue) of conventional digital cameras provide no insight into how a person will perceive a color. Thus, the measurement of a color error in this space does not readily translate into corrective action.

HUMAN-ORIENTED COLOR

To provide color feedback that has human significance, the vision system needs to use a color space that describes colors

only determining factor. The rate at which image processing can occur also faces limits. Several factors affect this rate, including image resolution, the type of processing needed, and the image processor's performance.

PROCESSING CHOICES

Most often, the application dictates the image resolution and type of processing, leaving designers only the choice of image processor to trade off cost and performance in the system design. Complete image-processing boards and systems in VME, PCI, and other board formats are available from companies such as Cognex, Dalsa, Epix, Matrox, National Instruments, and Philips Applied Technologies. In addition, some

cameras manufacturers incorporate image-processing hardware into their products, forming a "smart camera" that is user-programmable to handle many machine-vision tasks.

Some of the choices available at the component level for creating custom image-processing systems include general-purpose CPUs, DSPs, dedicated processing hardware, and specialty image pro-

Algorithm	Description	Usage
Histogram	Counts and graphs the total number of pixels at each gray-scale level	Determine whether the exposure and contrast of an image are suitable to the application
Spatial filters	Remove noise, sharpen, smooth, and transform images	Improve image quality
Blob analysis	Detect binary large objects (blobs) to find regions of interest in an image	Locate objects, detect flaws, identify objects
Edge detection	Locate the edges of objects	Measure objects, detect objects, determine alignment
Pattern matching	Locate regions of an image that match a template	Locate and identify objects for alignment, inspection, and measurement
Color matching	Compare the color of a region with a color template	Identification, sorting, color inspection




Figure A Vision systems that must measure colors and provide corrective feedback need to use a color space such as the CIE L-a-b, which describes color in human-perception terms (courtesy Edmund Optics).

in a manner similar to human perception. The CIE (International Commission on Illumination) L-a-b color standard is one such color space (Figure A). This standard uses "opponent theory," which relies on the fact that, in human perception, an object can look neither red and green at the same

time nor yellow and blue at the same time.

The CIE L-a-b system uses three values—Lightness (L), a^* , and b^* —to quantify a color. The first, L, quantifies perceived brightness (gray level). The second, a^* , represents how red or green the object looks. Positive a^* values

represent reddish colors, and negative values indicate greenish ones. The third parameter, b^* , indicates yellowish versus bluish colors. Positive b^* indicates yellow, and negative b^* indicates blue.

The CIE L-a-b color space is perceptually uniform; that is, equal differences in the color space represent equal human-perceived color differences. This representation simplifies the correction of color errors. If a measurement shows that a color looks too bluish, fixing the problem becomes simple.

A similar color space is the HCL (hue/chroma/lightness) standard. Hue describes how red, green, blue, or yellow the color appears. Chroma represents the color's departure from gray, which humans perceive as saturation or vividness. Lightness describes how dark or light an object is. The lightness scale runs from black to white with gray in the middle.

These two color standards map onto one another with CIE L-a-b values translating uniquely to the HCL color space. You can calculate hue and chroma from a^* and b^* values. Lightness is the same in either scale. Thus, a system that can work in one representation can work in both, providing information that will help evaluate and correct color errors.

The RGB color space also maps into the human-oriented color spaces but not completely. The mapping function to cover the entire color space describable by HCL or L-a-b requires negative values for red and green in some areas. Because cameras can provide only positive output values, working in the RGB space results in an inability to describe some colors. Applications that must handle the entire color spectrum require sensors such as spectrometers or colorimeters to make color measurements.




Figure 1 The image sensor in digital cameras is a CCD (charge-coupled device) that typically shifts out image data one pixel at a time, coupling image resolution with the maximum frame rate achievable (courtesy Imperx).

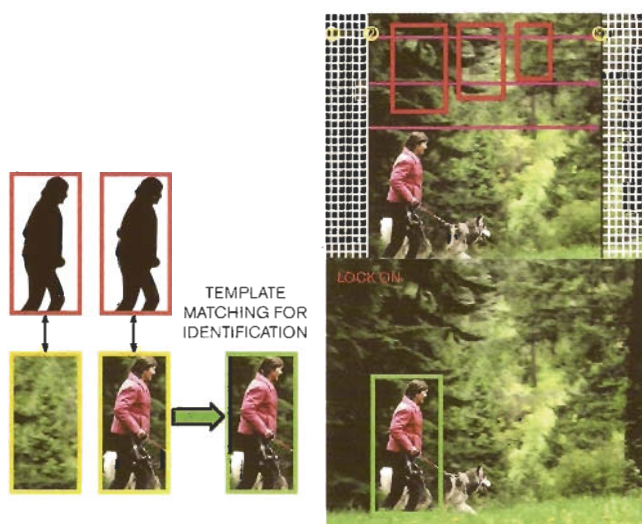


Figure 2 The latest wave of machine-vision-control systems is appearing in automobiles, identifying objects such as nearby pedestrians to warn drivers of impending hazards (courtesy NEC).

processors. General-purpose CPUs, including PCs, are most useful if the CPU is also slated for use in other system-control tasks and if the image-processing tasks are modest, with simple algorithms and relatively low resolution and low frame rate. Dedicated processing hardware—possibly implemented in an ASIC or an FPGA—is most effective at the other performance extreme: high resolution and high frame rate. The more complex the processing algorithm, however, the more difficult the hardware design.

For the highest performance with complex processing algorithms, a software-driven DSP or specialty image pro-

cessor may prove more cost-effective than dedicated hardware. A number of high-performance DSPs suitable for machine-vision applications are available from companies such as Analog Devices and Texas Instruments. These devices are well-established product lines with substantial development-tool and library support for machine-vision applications.

The specialty image processors represent a new breed of devices targeting the automotive market, although they are also useful for other vision-control applications. The automotive market for machine vision is growing, according to a

recent report (Reference 1). Among the automotive applications for machine-vision-control systems are those targeting automated parking and driver safety.

Automotive-safety applications use machine vision to supplement a driver's awareness of such things as lane markings, pedestrians, and other objects in the road by identifying and highlighting them in a video display (Figure 2). The vision systems may also monitor the driver, looking for indications of dozing, distraction, or inattentiveness. Currently, such systems provide warnings or alerts to the driver, such as by sounding an alarm or vibrating the steering wheel, rather than take control of the vehicle. Future uses may, however, also include automatic activation of braking systems or even evasive action.

The vision task for these systems is to scan images of the road or driver to locate, identify, and assess potential hazards. This task can be as simple as locating lane markings to determine whether the driver is drifting across the road or as complex as determining the presence of a pedestrian near the vehicle, measuring the relative positions and movements of the car and pedestrian, and determining whether a collision is imminent.

AUTOMOTIVE VISION

Because these specialty processors are handling safety-critical functions, they must operate in real time using complex algorithms to help minimize false alarms. Thus, they represent some of the most powerful programmable image processors available. One product in this category is the MobilEye EyeQ image processor, manufactured by ST-Microelectronics. The EyeQ uses two ARM RISC-processor cores and four proprietary VCEs (Vision Computing Engines) operating in parallel to provide high-performance scene recognition and interpretation in a single-chip device.

NEC Electronics has also announced a specialty image processor for automotive-machine-vision applications, the Imapcar for image recognition. This device uses 128 processing elements in parallel, using an SIMD (single-instruction-multiple-data) architecture to handle multiple parts of an image simultaneously. The company claims an equivalent processing power of 100G opera-

tions/sec with a 100-MHz clock.

Whether an application requires such processing power depends on the complexity of the image-processing algorithms as well as the image resolution and frame rate required. In addition, the efficiency of the algorithm's software implementation is a significant factor. Fortunately, libraries of common image-processing functions that are highly optimized for performance are widely available, both from processor manufacturers and independent software providers. Table 1 lists a number of important image-processing functions available in software along with their uses in machine-vision-control systems.

The availability of such software along with high-performance digital cameras and ever-increasing processing power has shifted machine vision away from being a specialty requiring high levels of technical skill and programming expertise. Machine-vision systems are becoming accessible to a wider variety of control applications, in which shape, size, color, and position of objects are the key factors determining how the system needs to react. EDN

REFERENCE

- 1 ABI Research, "Camera Based Automotive Systems," 2006.

AUTHOR'S BIOGRAPHY

Contributing Technical Editor Richard A. Quinell has been covering technology for more than 15 years after an equally long career as an embedded-system-design engineer.

FOR MORE INFORMATION

Analog Devices www.analog.com	MobilEye www.mobil-eye-vision.com
Cognex www.cognex.com	National Instruments www.ni.com
Cypress Semiconductor www.cypress.com	NEC www.necel.com
Dalsa www.corecoimaging.com	Philips Applied Technologies www.apptech.philips.com
Edmund Optics www.edmundoptics.com	STMicroelectronics www.st.com
Epix www.epixinc.com	Texas Instruments www.ti.com
Imperx www.imperx.com	Toshiba-Teli America www.toshiba-teli.com
Matrox Electronic Systems www.matrox.com	