Imaging for industry & science

Machine Vision Handbook

Contents

Introduction to UKIVA

The primary objective of the UK Industrial Vision Association is to promote the use of vision and imaging technology by industry and science in the UK. Its members include manufacturers and suppliers of vision systems and of specialised vision and imaging components, integrators of such systems, and other parties sharing the interest of promoting the use of machine vision.

For further information please contact:
UK Industrial Vision Association
PO Box 25, Royston, Herts SG8 6TL
Tel: 01763 261419 Fax: 01763 261961
Email: info@ukiva.org www.ukiva.org

Acknowledgement
The UKIVA wishes to acknowledge the assistance from its members in providing information, editorial input and diagrams, and to the Photonics Knowledge Transfer Network for financial assistance in updating and publishing this handbook.

The uses of machine vision

The applications of machine vision fall into three broad categories:

- Process control
- Quality control
- Non-industrial applications (for example, traffic control), which are outside the scope of this introduction, though of increasing importance.

Inspection by human beings often cannot keep pace with modern industry's requirements for speed of production and quality of product. People get tired and make mistakes, and the criteria that they apply during inspections are inevitably subjective. Also, in some cases, it is not humanly possible to perform the inspection task due to the environmental conditions.

The cameras and computer systems that make up a machine vision system, in contrast, carry out measurements with a constant precision at a pace that is set by the production process itself. These advantages have led to an increasing uptake of machine vision by industries around the world.

Worldwide, some of the applications of this technology have included:

- Inspection of the optical quality of screens for televisions and computers
- Inspection of the quality of paintwork during motor-car manufacture
- Inspection of bank notes during printing
- Checking electronic circuit boards
- Checking pharmaceutical packaging is properly filled
- Inspecting bottles to ensure they are properly filled
- Checking for flaws on tiles in the ceramics industry

Machine vision is concerned with the automatic interpretation of images of real scenes in order to control or monitor machines or industrial processes. The images may be visible light, but could also be of x-ray or infrared energy, and could even be derived from ultrasound information.

How does a vision system work?

In outline, a modern industrial vision system will consist of:

- An illumination system
  - Good illumination is especially important for taking images of products on a fast production line, but some applications may use ambient light
- A lens for the camera
  - Correct lens selection is important to achieve an optimal solution
- One or more cameras to acquire the images
  - Cameras can be analogue, but the price of high-specification digital cameras is falling so these are being adopted more widely
- An interface device to transfer the images to the processor/PC
- An image processor/PC
  - One option is to use 'smart cameras' that integrate image processing within the camera itself, negating the need to transfer images to an external PC
- An interface that highlights out-of-specification process or quality to the operator
  - This can be in the form of a software graphical interface or simple electronic signals to illuminate lights or operate a reject mechanism.

The incoming image – a continually varying two dimensional array of energy (i.e. light) levels – is divided into picture elements, known as 'pixels'. These form rows and columns covering the whole area of the image. A pixel can only have one energy level: in a monochrome image that will be a so-called grey level; in a colour image, the information describing the 'colour' of a pixel is more complex. The essential point is that a pixel cannot be subdivided into any smaller regions of different grey-level or colour. This process of pixelisation is a sort of spatial digitisation.

For each pixel, the energy-level information must also be digitised, i.e. the analogue (continuously varying) levels produced by the camera must be represented by a finite number of steps. In many applications it is sufficient to digitise a monochrome image to 8 bits per pixel, which equates to 256 'steps' to represent the grey level of each pixel. In more demanding applications it may be necessary to digitise to as much as 14 bits (or 16,384 levels). Colour images are more complex and can be represented in a variety of formats in addition to RGB (red, green, blue). Colour images typically contain two to three times more information than a monochrome image.

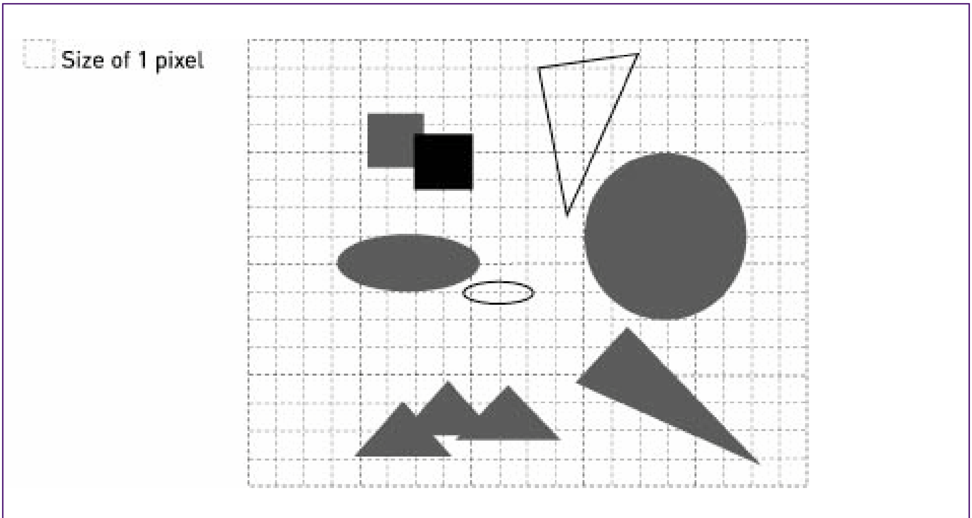

Figure 1

Some vision systems do not use an area array camera, but instead use a linear array camera, which produces just one line of 'image' at a time; a continuous two-dimensional image is generated as the product passes under the linear array, (typically on a conveyor belt) and this may be treated as a series of individual frames or the processing may be on a continuous basis to match the input.

How does the processor interpret the image?

The essential steps are segmentation and analysis – which are essentially software algorithms/functions that run on the processor.

Segmentation - this is deciding which parts of the image need interpretation/analysis and which are 'background'. Often it is possible and necessary to refine the segmentation. For instance if the application is to find cracks or scratches on glass, the first stage segmentation would typically find 'objects' which are cracks, scratches, dirt and dust. It is common to refine this result and 'remove' dirt and dust from the 'objects' to be analysed.

Analysis - Once the image has been satisfactorily segmented, the processor can make a variety of tests and measurements on the 'objects' of interest in the scene.

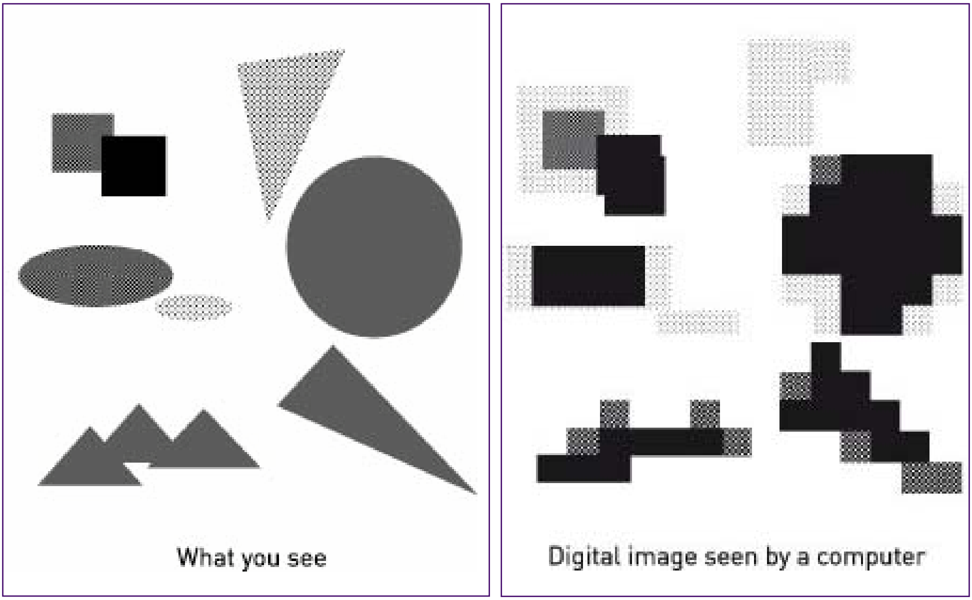

Figure 2

These are not real objects; they are parts of the image which represent objects of interest. The details of actual tests and measurements are beyond the scope of this handbook.

Once the image has been analysed, some form of output is required from the vision system, such as actuating a reject mechanism; updating statistical process control' information; and/or 'mapping' the location of defects and saving to file as a computer record.

What can a vision system do?

With increasingly more powerful processors the applications that can be successfully solved using vision systems are rapidly growing. An application may require one or more processing functions which when combined create a solution. The range of applications is vast and includes:

Inspection measurement

- Shape (or appearance) conformity checking
- Flaw detection – discrete items
- Flaw detection – web materials
- Colour confirmation
- One-dimensional measurement
- Two-dimensional measurement (within field of view accuracy)
- Structured light and other triangulation techniques
- Three-dimensional techniques (stereo, fringe methods, 'time of flight' techniques)

Recognition

- Character recognition
- Part (component or product) recognition

Guidance

- Predetermined guidance
- Continuous guidance

Special systems

- Systems designed for specific industrial applications.

Some of these applications by class are as follows:

Inspection – shape (or appearance) conformity checking

Current systems will generally start with a two-dimensional measurement operation to allow for any displacement of the object from the ideal position and will then carry out the 'pixel counting' checks for template matching or greyscale checking region of interest.

Inspection – Flaw detection (discrete items)

Flaw detection can be considered as a special case of conformity checking, in which the norm is 'no features'.

Examples include inspection of ceramic tiles at the 'biscuit' stage, before the glaze is applied or more specialised use of infrared inspection for flaws in glass bottles.

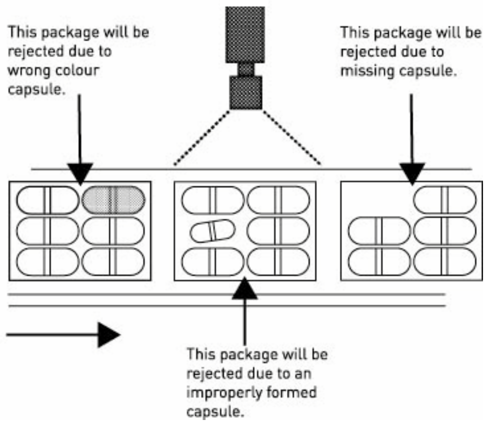

Figure 3

Inspection – Flaw detection (web materials)

The characteristic of many systems in this class is that they must examine large areas at high speed looking for relatively small flaws. Material produced in 'web' form includes sheet steel, paper, photographic film, glass, plastic sheet, and textiles. Some of these systems use mechanically-scanned laser beams rather than linescan cameras to form the image.

For wide webs an alternative is to use multiple linescan cameras, each covering only a part of the web.

Inspection – Colour confirmation

Colour inspection is widely used in the pharmaceutical industry to confirm that the correct coloured tablet or capsule has been placed in the right package

Measurement – One dimensional

One dimensional measurement can involve the measuring of the width of a continuously rolled or extruded product such as steel or rubber as it passes under the vision system. By rotating a sensor around the product, it is possible to deduce a two-dimensional cross-sectional profile from an essentially one-dimensional measurement provided that there are no concavities in the desired surface.

Measurement – Two dimensional (within field of view)

Two dimensional measurements can be made by superimposing 'optical callipers' on the image of the product. Measurement by this method is extremely fast because there is no mechanical movement, and the exact positioning of the object to be measured is not critical - the system can determine the X and Y displacement of the object and allow for the offsets and for a (usually limited) degree of rotation, transforming all co-ordinates to the actual, rather than theoretical, attitude of the object.

Some systems offer sub-pixel interpolation, which under appropriate conditions makes the system capable of measuring, and in particular detecting changes or differences in measurement, to one part in several thousand instead of the one part in several hundred which 'whole pixel only' systems offer. The number of pixels in the image can be increased by using very high resolution cameras, which are relatively expensive and may be rather slower to capture an image, or by using linear array (line scan) cameras and carefully measuring the rate at which the object moves under the line scan device to calibrate in the second dimension.

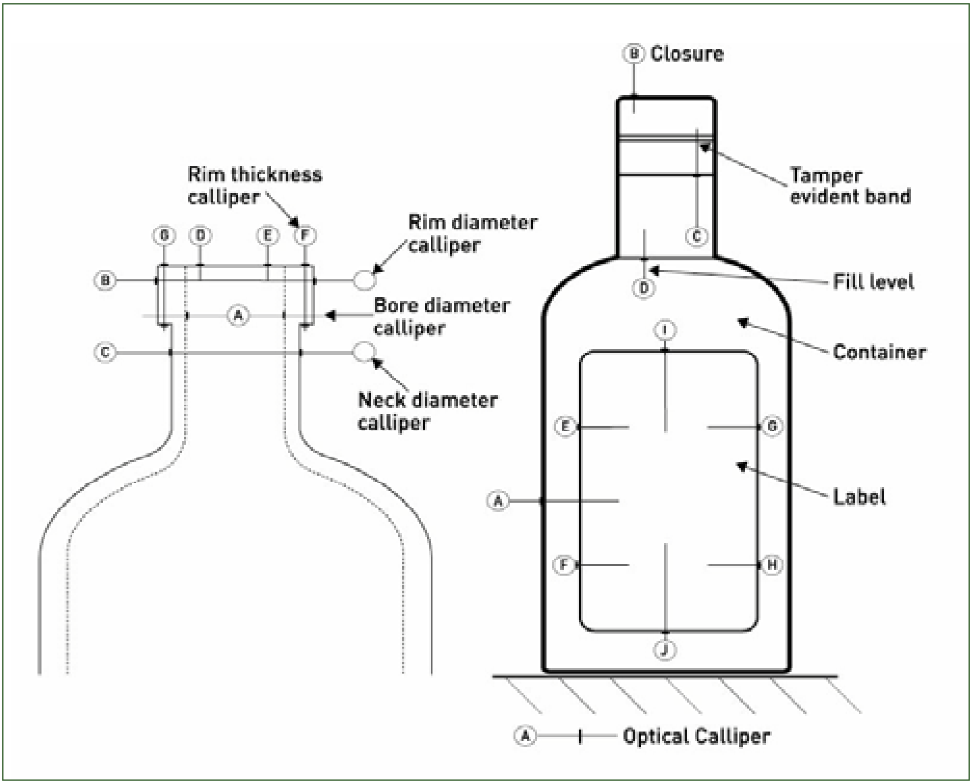

Figure 4: Measurement callipers

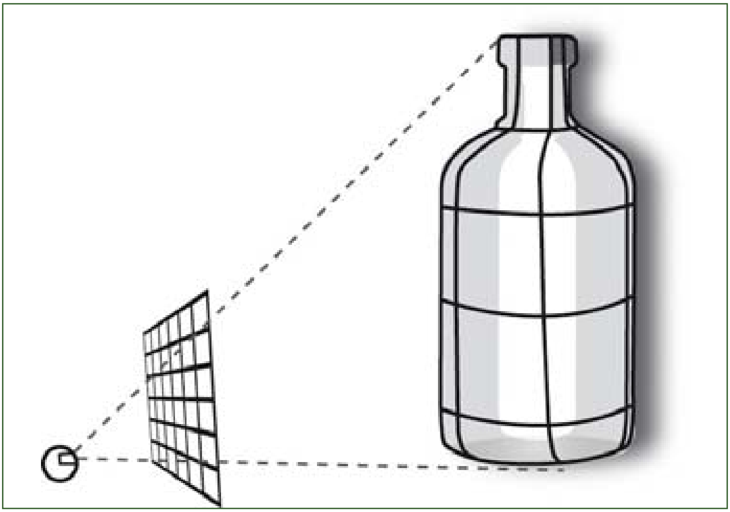

Figure 5

Measurement – 'Structured light' and other triangulation techniques

If a stripe of light is directed at a three-dimensional surface, and viewed by a two-dimensional camera with an angle between the optical axes of illuminating and viewing devices, the apparent shape of the stripe can be used to derive the shape of the surface along the length of the stripe. By scanning the stripe across the surface, an entire three-dimensional map of the surface can be deduced. It is often convenient to use a laser as the source of light, as it can easily be refracted to form a plane of light by passing it through a cylindrical lens, but any sharp 'edge' of light is equally useful. Developments in the use of programmable liquid crystal light valves allow three-dimensional shape information based on triangulation to be captured very rapidly.

Measurement – Three-dimensional 'full field' methods

Three-dimensional information can also be obtained using automatic stereo matching or fringe counting methods, or even 'time of flight' cameras which measure the time that light takes to reach each pixel from an illumination source from beside the camera.

Recognition – Characters

Character recognition can be based on simple correlation (comparison) techniques, without specialised knowledge of the way individual characters are formed or it can use more advanced algorithms that could, for instance, differentiate between a poorly formed '8' and an 'S'.

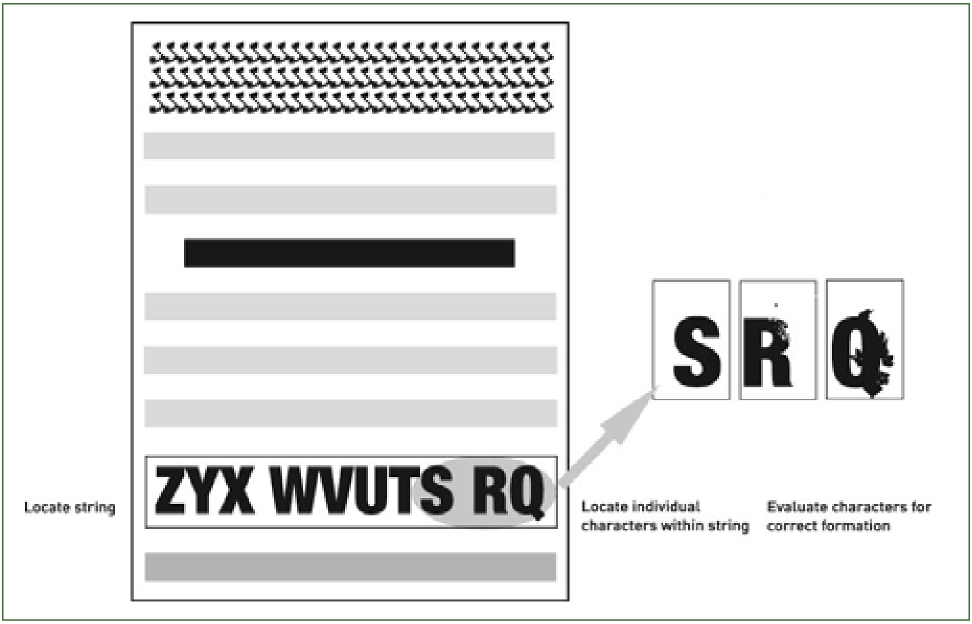

Figure 6

Vision systems can also be used to read bar-codes with a higher degree of confidence, and at longer ranges, than a laser scanner. Even if much of a barcode is obscured by mud or dirt, a vision system can usually read it whereas a conventional bar code reader will simply report it as unreadable.

Vision systems are also being used to read two-dimensional matrix codes that can contain far more information than a one-dimensional bar code.

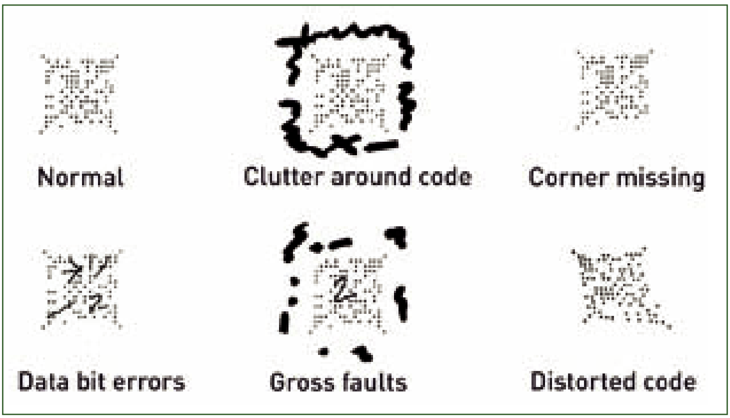

Figure 7: Matrix codes

Recognition – Parts and components

Part recognition can be used to identify which parts should be subjected to which process. For instance, a vision system can 'see' whether a car body is an estate, hatchback or saloon version, and instruct painting robots 'downstream' which of three patterns to follow.

Guidance – Predetermined

Pre-determined guidance is typified by a situation in which an overhead camera takes a 'snapshot' of the scene and the vision system directs a robot where to pick up or put down an object; the robot then works 'blind'. Typical tasks include de-palletising of heavy items such as engine blocks and crankshafts, but can also include much lighter work such as packing chocolates.

Guidance – Continuing

Continuing guidance most likely involves a camera mounted on the robot arm or 'hand' and the path of the robot is continually corrected by the vision system. A common application is arc-welding guidance, but the technique may also be used for controlling the path of robots spreading sealants or glues.

Special systems

All vision systems have some means of developing new applications, but if the platform on which the system has been designed is limited by the speed or current software functions, then specifically designed systems and software can often be developed to address new applications. Many UKIVA members offer considerable experience in customising systems to meet specific industrial needs.

Financial justification of machine vision

There are plenty of justifications for utilising machine vision:

Cost of materials

In most applications, the avoidance of making scrap products will probably result in a very short pay-off period.

To avoid making scrap, the automatic inspection system, whether providing 100% on-line inspection or being used off-line on samples taken from the line, often needs to be part of a statistical process control (SPC) system. This means that the system indicates whenever a parameter is drifting towards a tolerance limit, or is simply becoming unacceptably erratic, and corrective action can be taken before the limit is breached.

Cost of labour

The reduction in labour is also a significant cost saving, as many of the tasks performed by machine vision can directly replace people. The savings are even more significant for multi-shift operations.

In addition, the associated savings related to recruitment, benefits and no annual pay increases need to be considered.

However, the people involved in operating vision systems are often required to be more highly skilled.

Cost of quality

The increasing awareness of the cost of quality through such standards as IS09001 mean the use of machine vision can provide a more objective, reliable and consistent standard of inspecting products.

Savings in the optimisation of material usage, monitoring suppliers' quality and ensuring the quality of finished goods can lead to both tangible and intangible savings.

The cost of warranty repair work can be reduced, the avoidance or defence of litigation under product liability legislation can be important, and the improvement in customer confidence can lead to repeat orders and greater market share.

Interfacing with suppliers

Defining the system

Users want a problem solved which they have been used to expressing in their own terms. Suppliers of machine vision systems are able to specify what their systems can do in terms of an image having defined characteristics. Much of the art of concluding a mutually beneficial contract lies in ensuring that the gap between the three-dimensional, unconstrained, solid-object world of the user, and the two-dimensional, constrained-image world of the vision system supplier, is bridged in a way which each side understands and accepts.

It is important to realise that a specification which may be precise in the eyes of the prospective user may be very loose in image analysis terms. It may also impose much greater demands on the vision system than the user realises. The cost of a vision system tends to be linear with complexity and therefore it is important to have clearly defined requirements that have realistic boundaries.

Initial contact

Most suppliers of machine vision systems have to apply a selection process to incoming enquiries, selecting those enquiries which have obvious potential for profitable (and especially volume) business and a successful solution while treating the others (which may have hopelessly optimistic expectations of current vision technology) fairly cursively.

Here are some tips to help attract the suppliers' attention:

Problem

Outline requirements and have supportive details about vision applications, for example, product, size, type of fault, size of fault, line speed, number of lines, and working environment.

Justification

Assess the justification from the saving related to both tangible and in tangible benefits.

The justification can help determine whether an 'off the shelf' or 'customised' system will be the most appropriate.

Evaluation

Request a visit to site and have available a range of samples exhibiting the parameters to be inspected.

Occasionally, a feasibility study may be required and the cost of that may often be offset against a contracted system.

Specification supplier

Prepare a realistic specification that should address the key parameters. Until you actually start to use a vision system, you probably don't know what quantitative standards your inspectors really achieve, regardless of what their specification may say.

Concentrate on addressing the main problem and accept as a bonus the ability of the system to address secondary issues.

Contact a recognised supplier and review your problem with them. The UKIVA offers a database of members who have been supplying industrial inspection systems for a number of years. Many are specialists in particular industries and may be familiar with your type of product or process.

Ideally, this approach will yield one or more respondents who can demonstrate directly relevant experience.

In addition, other commercial consideration can be discussed, such as delivery, commissioning, training, documentation and warranty support.

Conclusion

In the last 15 to 20 years many thousands of machine vision systems have been installed worldwide in industrial applications.

New applications appear all the time as the impact of automation and demand for quality increase. For the first-time user, the savings can be significant, but understanding how they work and how they can be introduced is the key to their successful integration. The UKIVA hopes that this introduction will provide a guide to machine vision, a review of typical applications, the areas where savings can be made to justify their use and points to consider when approaching an established vision supplier.

Much more introductory material is freely available from the UKIVA website, www.ukiva.org, together with pointers to many other sources of comprehensive information about imaging and vision.

First published as 'Machine Vision Handbook' January 1997; revised April 2001; third edition Summer 2007

UK IVA

Imaging for industry & science