1

An introduction to computer vision

1.1 Computer vision: image processing or artificial intelligence?

What is computer vision and why would one be interested in studying it? It is perhaps easier to answer these two questions in reverse order. There are several reasons why one would be interested in computer vision, but the following two will serve to illustrate the many directions from which one can view the subject area:

1. All naturally occurring intelligent life-forms exhibit an ability to interact with and manipulate their environment in a coherent and stable manner. This interaction is facilitated by on-going intelligent interplay between perception and motion-control (i.e. action); visual perception is fundamentally important to most intelligent life.
2. Most manufacturers are concerned with the cosmetic integrity of their products; customers quite often equate quality of appearance with functional quality. So, to ensure the successful long-term marketing of an item, it is highly desirable that its appearance is checked visually before packaging and shipping. Likewise, it is desirable that the inspection process be automated and effected without human intervention.

These two motivations for the study of perception characterize two possible extremes of interest in the processing, analysis, and interpretation of visual imagery: from the philosophical and perhaps esoteric to the immediate and pragmatic. And the subject matter of everything between these two extremes presents one with wide and varied spectrums of commercial interest, difficulty and, indeed, success.

The answer to the first question (what is computer vision?) now becomes a little easier to identify. The world we live in and experience is filled with an endless variety of objects, animate and inanimate, and, to borrow a phrase from David Marr (of whom we shall hear more later in the book), it is by looking and seeing that we come to know what is where in this world. So, if vision is a means to an end – to know the world by looking – then computer vision is exactly the same except that the medium by which the knowledge is gained is now a computational instrument rather than the brain of some living creature. Without doubt, this is a very broad definition. But the subject matter of computer vision *is* this broad: topics such as image restoration, image enhancement, automated visual inspection, robot vision, computer-based image understanding of general three-dimensional scenes, and visual perception and cognition all fall under the umbrella of the term 'computer vision'.

Although for centuries man has been interested in solving the puzzle of how man comes to 'see', the first computational experiments in developing artificial machine vision systems were conducted in the late 1950s and, over the last twenty-five to thirty years computer-based vision systems of widely varying degrees of complexity have been used in many diverse areas such as office automation, medicine, remote sensing by satellite, and in both the industrial world and the military world. The applications have been many and varied, encompassing character recognition, blood cell analysis, automatic screening of chest X-rays, registration of nuclear medicine lung scans, computer-aided tomography (CAT), chromosome classification, land-use identification, traffic monitoring, automatic generation of cartographic projections, parts inspection for quality assurance industrial, part identification, automatic guidance of seam welders, and visual feedback for automatic assembly and repair. Military applications have included the tracking of moving objects, automatic navigation based on passive sensing, and target acquisition and range-finding.

As we have seen, computer vision is concerned with the physical structure of a three-dimensional world by the automatic analysis of images of that world. However, it is necessary to qualify the use of the word *image*. First, the image is a two-dimensional one and, hence, we inevitably lose information in the projection process, i.e. in passing from a three-dimensional world to a two-dimensional image. Quite often, it is the recovery of this lost information which forms the central problem in computer vision. Second, the images are digital images: they are discrete representations (i.e. they have distinct values at regularly sampled points) and they are quantized representations (i.e. each value is an integer value).

Computer vision includes many techniques which are useful in their own right, e.g. image processing (which is concerned with the transformation, encoding, and transmission of images) and pattern recognition (frequently the application of statistical decision theory to general patterns, of which visual patterns are but one instance). More significantly, however, computer vision includes techniques for the useful description of shape and of volume, for geometric modelling, and for so-called cognitive processing. Thus, though computer vision is certainly concerned with the processing of images, these images are only the raw material of a much broader science which, ultimately, endeavours to emulate the perceptual capabilities of man and, perhaps, to shed some light upon the manner by which he accomplishes his amazingly adaptive and robust interaction with his environment.

1.2 Industrial machine vision vs. image understanding

Computer vision, then, is an extremely broad discipline (or set of disciplines) and in order to get to grips with it, we need to identify some way of classifying different approaches. To begin with, we note that humans live and work within a general three-dimensional world, pursuing many goals and objectives in an unconstrained and constantly changing environment in which there are many varied and, often, ill-defined objects. Industrial automation, on the other hand, is given to performing single repeated tasks involving relatively few objectives, all of which are known and defined, in manufacturing environments which are normally constrained and engineered to simplify those tasks. Industrial systems do not yet work with *general* three-dimensional environments (although the environments they do work in are often much less structured than one would suppose) and vision systems for manufacturing still exploit many assumptions, which would not generally apply to unconstrained worlds with many objects and many goals, in order to facilitate processing and analysis. There is a considerable dichotomy between the two approaches – a situation which must change and is changing – it is for this reason that the final chapter is concerned with advanced techniques and their migration to the industrial environment. Let us look a little closer at each of these classes of computer vision.

Approaches associated with general environments are frequently referred to by the terms 'image understanding' or 'scene analysis'. The latter term is now quite dated as it typically refers to approaches and systems developed during the 1970s. Vision systems specifically intended for the industrial environment are often referred to generically as 'industrial machine vision systems'.

Image understanding vision systems are normally concerned with three-dimensional scenes, which are partially constrained, but viewed from one (and often several) unconstrained viewpoint. The illumination conditions may be known, e.g. the position of the room light might be assumed, but usually one will have to contend with shadows and occlusion, i.e. partially hidden objects. As such, the data or scene representation is truly a two-dimensional image representation of a three-dimensional scene, with high spatial resolutions (i.e. it is extremely detailed) and high grey-scale resolutions (i.e. it exhibits a large variation in grey-tone). Occasionally, colour information is incorporated but not nearly as often as it should be. Range data is sometimes explicitly available from active range-sensing devices, but a central theme of image understanding is the automatic extraction of both range data and local orientation information from several two-dimensional images using e.g., stereopsis, motion, shading, occlusion, texture gradients, or focusing. One of the significant aspects of image understanding is that it utilizes several redundant information representations (e.g. based on the object edges or boundaries, the disparity between objects in two stereo images, and the shading of the object's surface); and it also incorporates different levels of representation in order to organize the information being made explicit in the representation in an increasingly powerful and meaningful manner. For example, an image understanding system would endeavour to model the scene with some form of parameterized three-dimensional object models built from several low-level processes based on distinct visual cues. At present, image-understanding systems utilize both explicit knowledge (or models) and software-embedded knowledge for reasoning, that is, for controlling image analysis.

Most industrial machine vision systems contrast sharply with the above approach. The scenes in an industrial environment are usually assumed to be two-dimensional, comprising known isolated rigid parts, frequently with a contrasting visual backdrop. Lighting is almost always a critical factor and must be very carefully organized. Typically, the ambient room lighting will be totally inadequate, and even confusing, so that each inspection station will require its own set of dedicated lights, each designed for the task in hand. The images which industrial machine vision systems use are frequently two-dimensional binary images (pure black and white, with no intermediate grey-levels) of essentially two-dimensional scenes. There is normally just one simple internal object representation or model; the analysis strategy being to extract salient features (e.g. area, circularity, or some other measure of shape) and to make some decision, typically using feature-based discrimination. This process frequently uses software-embedded (hard-coded) knowledge of the scene.

There are two complementary areas of industrial machine vision: robot vision and automated visual inspection. Both of these use essentially the same techniques and approaches, although the visual inspection tasks are, in general, not as difficult as those involved in visual perception for robotic parts manipulation, identification, and assembly. This is because the inspection environment is usually easier to control and the accept/reject decisions required for inspection are often easier to determine than the location and identification information needed for assembly. The significant problems associated with robotic part handling, too, has meant that advanced three-dimensional robot vision has not received the attention it merits.

1.3 Sensory feedback for manufacturing systems: why vision?

The answer to this question must necessarily be double-barrelled:

1. We need feedback because the manufacturing system is not perfect and free of errors: we wish to ensure that we are informed when errors begin to creep into the process, so that we can take corrective action and ensure that quality and productivity are maintained.
2. We use vision because it is by far the most versatile sense available and conveys extremely rich information when compared with, e.g., sonar or infra-red sensing. Furthermore, unlike tactile sensing, it senses the environment in

a remote manner rather than having to be in contact with the objects being analysed.

Systems which are equipped with (useful) visual capabilities are inherently adaptive and can deal with uncertainty in the environment, or at least that is what one would hope for. The upshot of this is that, by incorporating vision in the manufacturing process, not only can we identify when things go wrong (e.g. in visual inspection) but the uncertain and variable nature of the manufacturing environment can be catered for.

Unfortunately, vision, while versatile, is also the most complex of the senses, due mainly to the fact that most information in visual images is implicitly coded and requires extensive processing and analysis to make it explicit. *Visual sensing is difficult*: in humans, ten of the estimated one hundred cortical areas in the brain are devoted to vision and much work remains to be done before we can claim to have even a modest grasp of visual sensing.

Given that one acknowledges that vision is (potentially) a very powerful sense, let us look at some of the motivations for using visual processes in the industrial workplace.

☐ *Safety and reliability*

Considerations of safety and reliability usually arise in environments which are hazardous for humans (e.g. in close proximity to a milling bit) or because manufactured parts are of critical importance and 100 per cent inspection is required (e.g. defects in a brake lining might conceivably cause loss of life).

Machine vision also facilitates consistency in inspection standards; such systems don't suffer from the 'Monday-morning' syndrome and their performance can be (and should be) quantitatively assessed.

☐ *Product quality*

High-volume production using humans seldom facilitates inspection of all parts but automated visual inspection techniques may make it feasible; this depends on the complexity of the task and the effective throughput that is required by the manufacturing system. The latter consideration is particularly important if the vision system is to be incorporated in an on-line manner, i.e. inspecting each part as it is manufactured.

☐ *Flexible automation*

In environments where quality assurance is performed by a machine, it is feasible to integrate the inspection task into the complete production or manufacturing cycle, and allow it to provide feedback to facilitate on-line control. This provides for the adaptive requirements of AMT (advanced manufacturing technology) systems and facilitates the overall control by computer, such as is found (or, more realistically, as will be found in the future) in advanced computer integrated manufacturing (CIM) environments.

Integration within a CIM system is, however, one of the problems that is forestalling the rapid deployment of vision technology in advanced manufacturing environments, not least because the vision system will not usually know how to communicate with, e.g., the CAD (computer aided design) database or another manufacturer's robotic palletizing system.

1.4 Examples of industrial machine vision problems and solutions

Of the two sub-sections of industrial machine vision, automated visual inspection is presently by far the most important area, accounting for at least 80 per cent of all current applications. Within inspection, there is a great diversity of uses in fields such as quality assurance, machine monitoring, and test and calibration. These applications may be classified on a functional basis as either gauging, inspection, or sorting:

- Gauging is concerned with the measurement of dimensional characteristics of parts and with checking tolerances.
- Inspection, *per se*, is concerned with performing part verification, i.e. establishing whether there are any parts or sections missing from the object being inspected or whether there are any extraneous parts which should not be present. Alternatively, one might be interested in performing flaw detection, i.e. effecting the *detection* and *classification* of flaws (usually surface flaws) on the part: for example, the detection of scratches on plastic housings.
- Sorting is concerned with the identification and recognition of parts. Parts will usually be on a conveyer system and it is pertinent to note that this does not necessarily involve robot manipulation as the part can quite simply be pushed into an appropriate bin using a simple pneumatically actuated flipper.

The applications of robot vision are less varied. At present, this is primarily due to the uncertain nature of the industrial environment within which a robot will be operating: there are typically three dimensions to deal with instead of two, partly manufactured objects tend to be poorly finished having spurious material attached (e.g. swarf) and they tend not to appear exactly where they are expected or are supposed to be. The two application areas which are well developed are materials handling and welding. The reasons for this are that these problems can often be reduced to two-dimensions; objects on a conveyer or pallet are visualized as two-dimensional silhouetted shapes and tracking a welding seam is conceived, locally, as a two-dimensional exercise in the estimation of the disparity between the position of the welding rod and the metal seam. The application which one would expect to be the main subject matter of robot vision – part assembly – is very poorly developed for the reasons given above and, additionally, because the assembly operation requires greater dexterity on the part of the robot than is currently

offered by commercial manipulators. It is for this reason that the development of compliant manipulation and the incorporation of tactile sensors in the robot grippers are fundamentally important.

Despite these negative comments, there is still a great deal of potential for existing vision technologies; there are many common industrial problems which can be successfully and cost-effectively solved with the judicious use of well-understood and simple vision tools. To illustrate this and some of the techniques that have been alluded to so far, a few very simple examples are in order; a more sophisticated application is described in detail in Chapter 8.

1.4.1 Measurement of steel bars

Raw material for steel ingots to be used in a casting process is delivered in four-foot long bars with a circular cross-section. Individual ingots are then cut from the bar and stacked in a wooden box before being moved to the next stage of the casting process. To ensure high-quality casting with minimal wastage of raw material, it is essential that the weight of the charges fall within acceptable tolerances. The nominal tolerances on weight are ± 5 g.

If the cross-sectional area of the bar of raw material were constant, the required weight would be given by an ingot of fixed length. Unfortunately, the cross-sectional area is not constant and the outside diameter of the bar can vary considerably (a 1.25″ bar has a nominal variation of 0.01″ but often exceeds this limit). Thus, in order to ensure the correct weight of a cut charge, the cross-sectional area of the bar must be monitored and a length of bar chosen such that the volume (and hence weight) falls within the required tolerance.

To measure the cross-sectional area, an 'outside diameter' gauge comprising a diametrically opposed light source and an image sensor is mounted on a robot end effector and moved along the axis of the bar (see Figure 1.1). In this manner, the

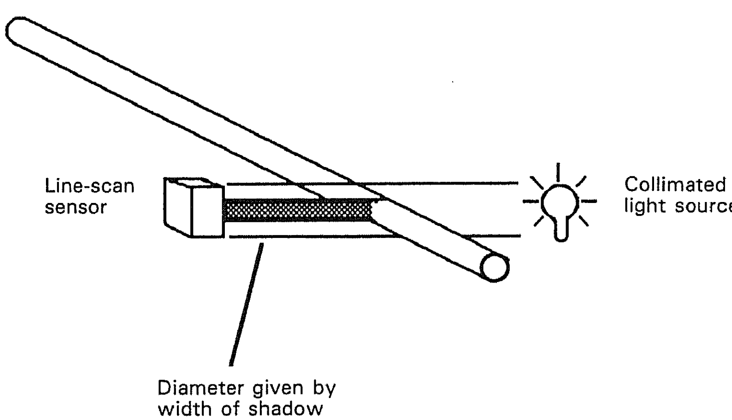

Figure 1.1 Measurement and inspection of steel bars.

bar will form a 'shadow' on the light sensor, the size of which corresponds to the diameter of the bar at that point. Thus, a profile of the bar can be built as the bar passes through the gauge. One pass of the sensor generates a single signature describing how the diameter varies over the scanned length. If it is assumed that the bar is truly circular in cross-section, this single measure is sufficient to compute the volume of the bar; several passes of the sensor at different orientations (i.e. rotating the sensor about the axis of the bar after each pass) allow this assumption to be dropped. The essential point to be noticed about this system is that the techniques used are inherently simple (measuring the length of a shadow cast directly on a sensor) and do not require any special sophisticated software. The gauging process is very fast and consequently can be integrated into the production system as a feedback device to control the ingot cutting operation directly.

1.4.2 Inspection of computer screens

Cathode ray tubes (CRTs) must be inspected and calibrated during the assembly of personal computers. The picture position, width, height, distortion, brightness, and focus need to be computed and, if any of these measurements lie outside specified tolerances, real-time feedback is required to facilitate on-line adjustment. Such a system has been configured using standard off-the-shelf equipment, comprising an

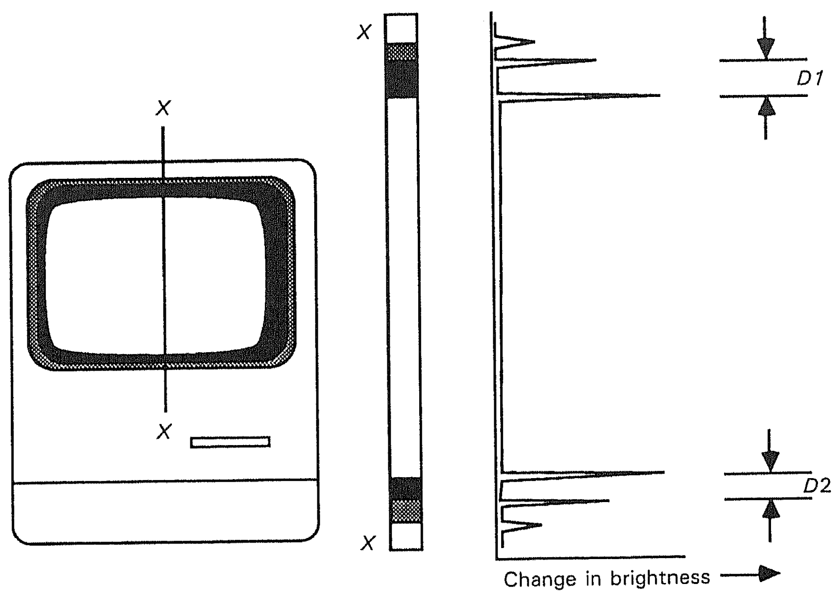

Vertical eccentricity $= D1 - D2$

Figure 1.2 Inspection of computer screens.

Apple Macintosh II, a data translation image acquisition device (framestore), and a Panasonic CCD (charge coupled device) camera.

For the purposes of illustrating the image analysis, consider the problem of identifying the position of the picture on the screen, such as is depicted in Figure 1.2. We wish to measure the distance between the edge of the picture (shown in white) and the inside edge of the computer face-plate, i.e. the width of the black border. To do this, we sample a small linear section (shown as the vertical *X–X* section in Figure 1.2) and identify the transitions between white and black and between black and grey. This is accomplished by estimating the rate of change in intensity or colour along the strip. A signature of this intensity change is shown in Figure 1.2; the required distances, *D*1 and *D*2, correspond to the distance between the first two peaks we encounter as we journey from the centre of the strip to the periphery. The CRT is adjusted until these two distances are, to within specified tolerances, equal. All the other features which need to be checked, with the exception of focus and brightness, are computed in a similar manner.

1.5 A typical system architecture

In the context of a manufacturing environment, a machine vision process will typically encompass nine components: the manufacturing assembly line; some mechanism for delivering the parts to the machine vision system; a method of automatically presenting the part to the system (including part acquisition, part positioning, and registration); some sensing system which provides a representation understandable by computer systems (and digital systems, in general); the representation itself; some process to extract pertinent features; a set of criteria upon which some decision process can be based; a method facilitating the automatic release of the system; some mechanism for collecting the processed part (see Figure 1.3).

Rather than discuss in detail each component of the imaging system, as this would pre-empt the subject matter of later chapters, we will present here no more than a schematic overview (illustrated in Figure 1.4) denoting many components and terms which have not yet been explained; the reader is asked to bear in mind that their introduction here is by way of preview rather than a self-contained exposition.

The imaging part of the machine vision system comprises sub-systems for image formation, image acquisition and (pre-)processing, image analysis, and image interpretation. The image formation system, consisting of the part illumination, the sensing element and the associated optics, is critical to the successful deployment of the system. This system generates an analogue representation of the image scene, typically in the form of a conventional TV video signal. The task of the image acquisition and processing sub-system is to convert this signal into a digital image, and to manipulate the resultant image to facilitate the subsequent extraction of information. The image analysis phase, working with

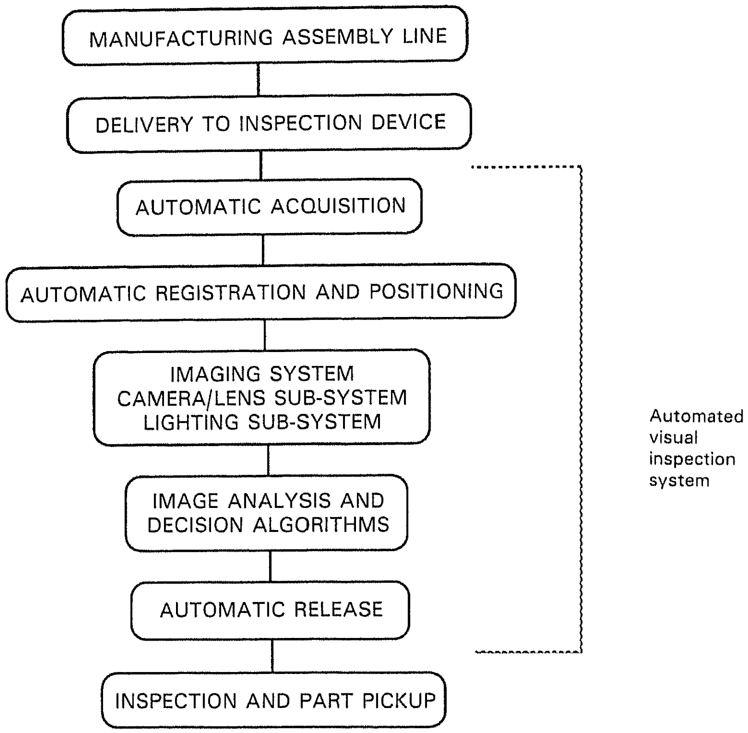

Figure 1.3 Components of a machine vision system in the context of a manufacturing environment.

the processed image, is concerned with the extraction of explicit information regarding the contents of the image (e.g. object position, size, orientation). Thus, there is a fundamental difference between image processing and image analysis: the former facilitates transformations of images to (hopefully, more useful) images, while image analysis facilitates the transformation from an image to an explicit (symbolic) piece of information. The final phase, image interpretation, is concerned with making some decision (e.g. accept or reject) based upon the description of what is in the image; it provides the link back to the environment to control the inspection process or the robot manipulator.

Each sub-system requires its own specialized type of hardware: image formation uses camera technology, optics, and various light sources. Image acquisition requires an analogue-to-digital converter or 'frame-grabber' to capture a frame of video information; image processing is a computationally intensive task (due to the amount of data comprising an image) and the predominant trend is to accomplish as much of this in hardware as possible. Image analysis typically takes place on a conventional micro/mini computer, though, again, there is a trend to accomplish as much of the analysis as possible in dedicated hardware. Image

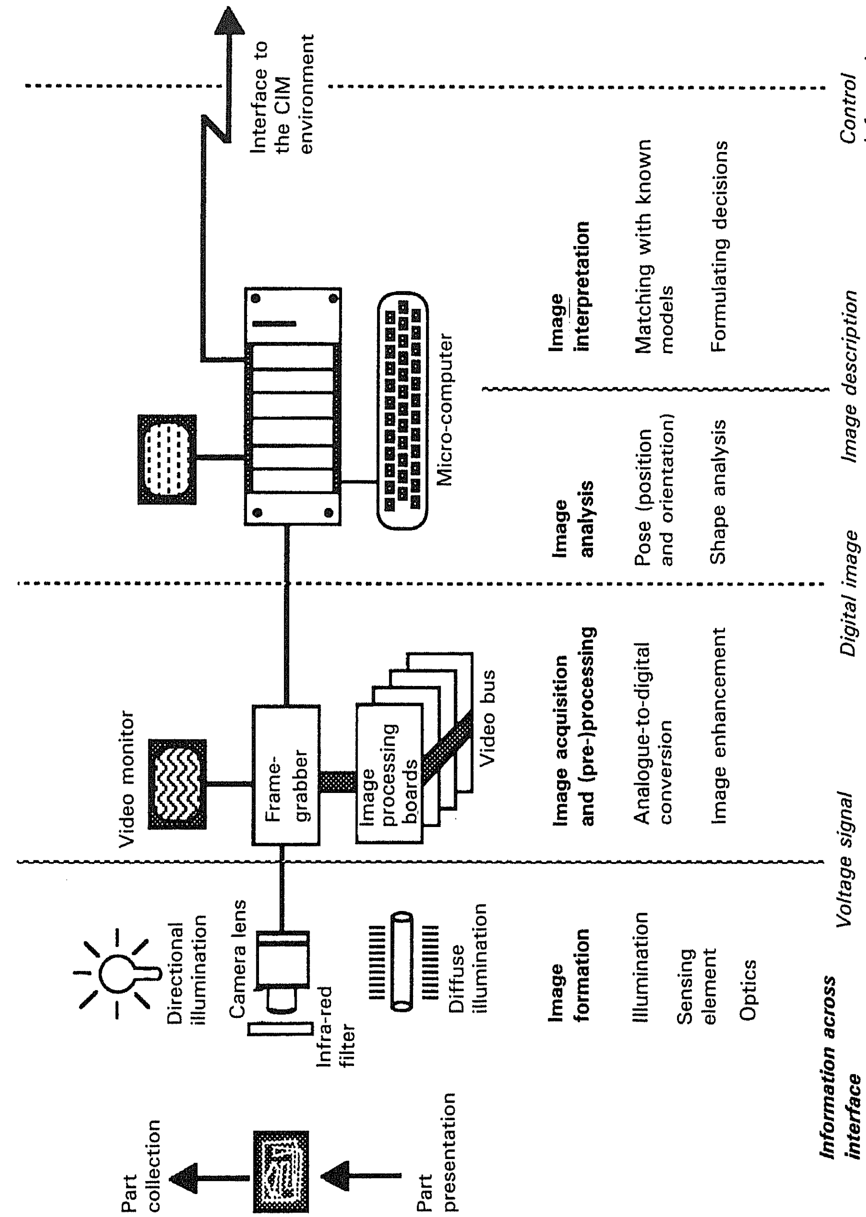

Figure 1.4 Typical system architecture (after Sanderson 1983).

interpretation is normally effected in software on a conventional computer, usually the same one which implements the image analysis.

Exercises

1. Comment on the validity of the statement that 'Industrial machine vision and image understanding have nothing in common'.
2. Given that an inspection task typically comprises three main stages of sensing, image processing, and image analysis, identify the component processes within each stage by describing the functional structure of a task involving the automated visual inspection of O-ring seals. The objective of the inspections task is to compute the outside and inside diameters and the diameter ratios and to check that each of the three values is within a certain tolerance. Pay particular attention to the components related to image analysis.
3. Describe briefly the essential difference between information contained in images and the information required to make a decision in a manufacturing environment.

References and further reading

Ballard, D.H. and Brown, C.M. 1982 *Computer Vision*, Prentice Hall, New Jersey.

Batchelor, B.G., Hill, D.A. and Hodgson, D.C. (eds) 1985 *Automated Visual Inspection*, IFS (Publications), Ltd, United Kingdom.

Besl, P.J., Delp, E.J. and Jain, R. 1985 'Automatic visual solder joint inspection', *IEEE Journal of Robotics and Automation*, Vol. RA-1, No. 1, pp. 42–56.

Binford, T.O. 1982 'Survey of model-based image analysis systems', *The International Journal of Robotics Research*, Vol. 1, No. 1, pp. 18–64.

Brady, J.M. (ed.) 1981 *Computer Vision*, Elsevier Science Publishers, The Netherlands.

Brady, M. 1982 'Computational approaches to image understanding', *ACM Computing Surveys*, Vol. 14, No. 1, pp. 3–71.

Bolles, R.C. 1981 *An Overview of Applications of Image Understanding to Industrial Automation*, SRI International, Technical Note No. 242.

Boyle, R.D. and Thomas, R.C. 1988 *Computer Vision – A First Course*, Blackwell Scientific Publications, Oxford.

Chin, R.T. and Harlow, C.A. 1982 'Automated Visual Inspection: A survey', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-4, No. 6, pp. 557–73.

Christiansen, D. 1984 'The automatic factory', *IEEE Spectrum*, Vol. 20, No. 5, p. 33.

Connors, R.W., McMillin, C.W., Lin, K. and Vasquez-Espinosa, R.E. 1983 'Identifying and locating surface defects in wood: part of an automated lumber processing system', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-5, No. 6, pp. 573–83.

Corby, N.R. 1983 'Machine vision for robotics', *IEEE Transactions on Industrial Electronics*, Vol. IE-30, No. 3, pp. 282–91.

Caudrado, J. L. and Caudrado, C. Y. 1986 'A.I. in computer vision', *Byte*, Vol. 11, No. 1, pp. 237–58.

Duda, R.O. and Hart, P.E. 1973 *Pattern Classification and Scene Analysis*, Wiley, New York.

Fairhurst, M.C. 1988 *Computer Vision for Robotic Systems*, Prentice Hall International (UK), Hertfordshire.

Fischler, M.A. 1981 *Computational Structures of Machine Perception*, SRI International, Technical Note No. 233.

Gonzalez, R.C. and Safabakhsh, R. 1982 'Computer vision techniques for industrial applications and robot control', *Computer* Vol. 15, No. 12, pp. 17–32.

Gonzalez, R.C. and Wintz, P. 1977 *Digital Image Processing*, Addison-Wesley, Reading, Mass.

Hall, E.L. 1979 *Computer Image Processing and Recognition*, Academic Press, New York.

Hanson, A.R. and Riseman, E.M. 1978 'Segmentation of natural scenes', in *Computer Vision Systems*, Hanson, A.R. and Riseman, E.M. (eds), Academic Press, New York.

Hara, Y., Akiyama, N. and Karasaki, K. 1983 'Automatic inspection system for printed circuit boards', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-5, No. 6, pp. 623–30.

Horn, B.K.P. 1986 *Robot Vision*, The MIT Press, Cambridge, Massachusetts.

Jarvis, J. 1982 *Computer Vision Experiments on Images of Wire-Wrap Circuit Boards*, Conference Record on 1982 Workshop on Industrial Applications of Machine Vision, pp. 144–50.

Kelley, R.B., Birk, J. R., Martins, H.A.S. and Tella, R. 1982 'A robot system which acquires cylindrical workpieces from bins', *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-12, No. 2, pp. 204–13.

Kelley, R.B., Martins, H.A.S., Birk, J.R. and Dessimoz, J-D. 1983 'Three vision for acquiring workpieces from bins', *Proceedings of the IEEE*, Vol. 71, No. 7, pp. 803–21.

Mahon, J., Harris, N. and Vernon, D. 1989 'Automated visual inspection of solder paste deposition on surface mount technology printed circuit boards, accepted for publication in *Computers in Industry*.

Moore, F.W. 1987 'Remote visual inspection of nuclear fuel pellets with fibre optics and video image processing', *Optical Engineering*, Vol. 26, No. 2, pp. 152–5.

Nevatia, R. 1982 *Machine Perception*, Prentice Hall, New Jersey.

Pau, L.F. 1983 'Integrated testing and algorithms for visual inspection of integrated circuits', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-5, No. 6, pp. 602–8.

Pau, L.F. 1984 'Approaches to industrial image processing and their limitations', *Electronics and Power*, February, pp. 135–40.

Pinker, S. (ed.) 1984 'Cognition', *The International Journal of Cognitive Psychology*, Vol. 18, Nos. 1–3.

Pratt, W.K. 1978 *Digital Image Processing*, Wiley, New York.

Pugh, A. (ed.) 1983 *Robot Vision*, IFS (Publications) Ltd, United Kingdom.

Rosenfeld, A. 1984 *Why Computers Can't See (Yet)*, Abacus, Vol. 1, No. 1, pp. 17–26, Springer Verlag, New York.

Rosenfeld, A. and Kak, A. 1976 *Digital Picture Processing*, Academic Press, New York.

Rosenfeld, A. and Kak, A. 1982 *Digital Picture Processing*, Academic Press, New York.

Sanderson, R.J. 1983 *Machine Vision Systems: A Summary and Forecast*, Tech Tran Corp. Illinois.

Spoehr, K.T. and Lehmkuhle, S.W. 1982 *Visual Information Processing*, Freeman and Co., San Francisco.

Suresh, B., Fundakowski, R.A., Levitt T.S. and Overland, J.E. 1983 'A real-time automated visual inspection system for hot steel slabs', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-5, No. 6, pp. 563–72.

Vernon, D. (Rapporteur) 1987 *Sensoric Feedback for Control and Decision Support Systems*, Report of a Workshop in the Framework of ESPRIT Conference 1987, Directorate General XIII (eds.).

Vernon, D. 1989 'Computers See the Light', *Technology Ireland*, Vol. 21, No. 1, pp. 21–3.

2

Illumination and sensors

2.1 Illumination

Scene and object illumination play a key role in the machine vision process. The central purpose of imposing controlled constant illumination is to enhance visually the parts to be imaged so that their flaws, defects, and features are highlighted and so that their identification and classification by the vision system becomes somewhat easier. Although the choice of lighting will typically be application-dependent, some general points may be pertinent.

The common incandescent bulb is probably the simplest source of light. It is cost-effective and it is easily adjusted for light intensity; however, it generally provides directional illumination since it is, to an approximation, a point source of light. Hence, incandescent bulbs cast strong shadows which invariably cause problems for machine vision software. Special bulbs are normally required as degradation in emitted light intensity is common with age. Furthermore, incandescent bulbs emit considerable infra-red radiation; this does not cause problems for humans as we are not sensitive to such light but some camera sensors, particularly so-called CCD cameras, are sensitive and visual data can be washed out by the reflected infra-red rays.

For most machine vision applications, a diffuse source of light is the most suitable. Diffuse lighting is non-directional and produces a minimum amount of shadow. Fluorescent lighting is the simplest and most common method of obtaining diffuse illumination and is especially good for providing illumination of large areas.

In situations where the only features that need to be inspected are evident from the silhouette of the object, back-lighting is the most appropriate. Back-lighting, e.g. in the form of a light table, provides high contrast between the object and the background upon which the object rests. Its advantage is that it facilitates very simple object isolation or *segmentation*, a topic to which we will return in Chapter 5.