

- [Scientific Computing World](#)
- [Research Information](#)
- [Laser Systems Europe](#)
- [Imaging & Machine Vision Europe](#)
- [HPC Projects](#)
- [Europa Science](#)
- [Electro Optics](#)



- [HOME](#)
- [NEWS](#)
- [PRODUCTS](#)
- [SUPPLIERS](#)
- [EVENTS](#)
- [FEATURES](#)
  
- [FEATURES INDEX](#)

## 3D VISION

# In stereo

Stephen Mounsey looks into the depths of the z-axis to report on 3D imaging techniques

### [Imaging and Machine Vision Europe: June/July 2009](#)

In the movie 300, King Leonidas asks his captain 'I trust that "scratch" hasn't made you useless?' to which the captain replies 'Hardly, my lord; it's just an eye. The gods saw fit to grace me with a spare.' While nobody born outside of ancient Sparta would be quite so blasé about the loss of a sense organ, it is true that we can still understand our 3D world using only one eye. This is also the case with machine vision; 3D measurements of the world can be taken with either one camera or several.

When observing a visual field, our brains are able to perceive depth in several ways. For most people, the effect of binocular disparity plays a key role – the brain uses parallax differences between the images seen by each eye to form a 3D model of the world. In addition, objects occluding one another, knowledge of shadows, perspective based on the known size of an object, and the way our view changes as we move our heads can all



A ZPS-1000 laser scanning 3D measurement system from Z-Laser. The specialised camera and scanning lasers are held at a fixed angle, and objects to be imaged are moved beneath the system (typically by a conveyor). Image courtesy of Alrad Instruments.

add depth clues to an otherwise 2D scene. Indeed, surveys show that as many as 20 per cent of us rely solely on visual depth clues and forego binocular disparity completely.

With respect to 3D machine vision, there are two main groups of solutions using either one camera or several. Single-camera solutions use similar depth clues to those our brains use, but seeing as how computers are not as well adapted to the task as human brains, they often need the help of a scanning laser beam. The laser beam serves as a reference point, providing information about where the viewpoint is in relation to the object. Multi-camera solutions tend to use stereoscopic effects exclusively. Both approaches have two requirements: calibration and conformation.

Conformation means identifying a feature in two or more images, whether they're taken simultaneously by two cameras, or sequentially by one camera. This can involve either feature recognition algorithms or, more commonly, some kind of added reference point (a laser line, or a high-contrast shape). Calibration involves teaching the system where its various components are in relation to each other, and, in some cases, teaching it where the objects it is looking at will be.

### **Single-camera vision - laser scanning**


In industrial applications, by far the most common technique for taking 3D measurements of a subject involves laser scanning. Here, a camera is trained on a laser spot as it scans across a test subject. Analysing the relative position of the laser reflection in the camera's field of view gives a value of the elevation of the spot at any given time. Turn the spot into a line, and the technique yields a cross-sectional profile of the subject. In this technique, the angle between the laser plane and the camera must be fixed, and so it is usually the subject that is moved, while the laser and camera are held stationary. Many applications are related to quality control on a production line, with the camera-laser system held at a fixed height.

In laser scanning applications, frame rate becomes one of the limiting factors. Ian Alderton, technical sales director of Alrad, explains: 'To get higher 3D resolution, you need to take more pictures, and take them closer together. Typically, a 3D imaging system might view 1,000 x 1,000 pixels at a rate of 500 frames per second. This is too much information to get through a standard Camera Link or GigE adapter in order to do the analysis on a PC.' Where bandwidth between the PC and the camera is a limiting factor, the scanning process can be accelerated (or given higher resolution) by way of on-board processing within the camera. By carrying out the line-recognition parts of the process on the camera side, only the positional data need be transmitted to the PC for further processing. 'The camera is able to throw away nearly 80 per cent of the data, enabling higher framer rates,' says Alderton. This on-board pre-processing can be achieved either with specialist 3D cameras or by way of off-the-shelf smart cameras, programmed in the correct way and coupled with the correct software.

Arnaud Lina, Matrox Imaging Library software leader at Matrox Imaging, explains that calibration of the system to all three dimensions, x, y, and z, is not a trivial matter: 'One of the dimensions comes from the conveyer, and is purely mechanical. A second component arises from the field of view of the camera, and the third arises from the parallax effect of the laser plane relative to the camera. Different phenomena lead to different sources of inaccuracies in calibration, and all of these inaccuracies need to be taken account of in order to resolve x, y, and z positions in real space. These problems are all addressed on the software side of things; we find ways to calibrate the distances involved. The other consideration is making each variable as easy as possible for the user.'

Ultimately, 'as easy as possible' will mean that a newly set-up imaging system can be calibrated in a single step, using a calibration object of known dimensions and geometry. Currently, calibration involves two steps, and can take a couple of hours: first the camera is calibrated against a grid, and secondly the laser position is established with respect to the camera.

High throughput is important in many quality control applications, but a 3D solution only becomes necessary in certain circumstances. For example, one automotive client is known to use a laser-scanning approach in order to read the raised writing on the rim of car tires – black on black would not be visible using any other technique. There are, however, many circumstances that demand a 3D solution but where high throughput and high resolution are less of a priority. In these cases, the expense of a high frame rate, preprocessing camera can be avoided by using a more standard machine vision camera, operating at around 25fps. Matrox's Lina describes these simpler solutions: 'A lot of the time, the imaging system is capable of more speed than is actually needed for an application. Because these [high-speed] systems are dedicated niche devices, they tend to be expensive. At Matrox we addressed this by putting together a software-based system, with which users are able to build a 3D imaging system from off-the-shelf components.'



*3D systems based on laser scanning can make use of specialist high frame rate cameras for high throughput and high resolution applications. In order to meet bandwidth constraints, recognition of the scanning laser line is carried out on-board the camera. Image courtesy of Alrad Instruments.*

In situations where speed of the imaging system is limited by the processing speed of the PC, rather than being limited by the bandwidth between the camera and the PC, further versatility can be added to the laser scanning approach by way of specialised processing boards mounted in PCs. These expansion cards carry out a majority of the 3D processing, leaving the PC itself free to carry out other tasks.

The technique of laser scanned 3D imaging is not without its drawbacks. Firstly, there are several limitations in terms of the kind of object that can be scanned. In order for a laser spot to be visible to the observing camera, the material of the object must be reflective to the wavelength of light used. Lasers used are usually either 635nm in the red or 532nm in the green parts of the visual spectrum. Black objects will not reflect these colours adequately. Furthermore, due to the resolution limits of the camera, there exists a size/accuracy trade-off, in terms of the object and the 3D resolution respectively. 'It is not possible to achieve accuracy and size [of the object] at the same time. The bigger your object is, the more difficult it is to have a high accuracy; the camera will need to be further away, and the laser will be at a higher angle,' explains Lina. For complex object geometries, occlusion may also arise if the laser line is hidden from the camera's view at any point. Occlusion can be overcome by adding more lasers, but this increases the complexity of calibration, and reduces throughput.

## **Stereoscopic vision**

Laser-scanning is not suitable to many applications, whether because of a lack of fixed geometry, unknown object size, or lack of movement in the system. Stereoscopic vision systems, using two or more cameras, can be used to create a more human-like vision system, comparing the images captured by each camera in order to create a 3D image based on parallax effects. The technique relies heavily on conformation, i.e. being able to pick out features in two images and identify them as the same object. While cameras spaced far apart will see a larger parallax effect, and measure

depth more accurately, the difficulty of identifying features are compounded by the fact that the cameras are viewing the subject from different angles. Many solutions use an array of cameras.

Stereoscopic imaging is used in a more diverse range of applications than laser scanning. Robotics is a field in which stereoscopic imaging is employed to give a machine a view of its surroundings. Bin picking tasks, wherein a system must automatically choose the correct object, are solved by way of stereoscopic imaging systems.

3D imaging is rarely of importance to consumers, save for one area – video games. Motion capture has become an important part of game design, as studios seek to put ever-more lifelike characters into increasingly choreographed situations. A typical motion capture setup consists of actors in black suits, with ping pong balls sewn on at key points on their body as reference points. Two or more cameras are trained on the actor, and the resulting wireframe is subsequently textured.

## Time of flight

Microsoft has recently announced something of a next-generation peripheral for its Xbox 360 console, named Natal. The system is essentially a 3D imaging camera, although it uses a principle which differs from either laser scanning or stereoscopic imaging. The sensor consists of an RGB camera, an infrared illuminator, and a monochromatic CMOS camera. Depth is established using a time-of-flight approach, in which the time difference between a pulse of infrared being emitted and being detected is compared to the speed of light to give an accurate estimate of the distance to the subject.

Although the applications to which 3D imaging is applied are currently rather specialist, it is clear that novel uses for the technologies continue to appear. Despite the increased complexity of current 3D systems relative to conventional machine vision, it seems that the market for 3D vision will continue to increase, both in industrial applications and in commercial gadgetry.

## Search features

## Subscribe

*Imaging and Machine Vision Europe* is circulated six times a year at no cost to readers with suitable qualifications; put simply, you will be able to get the magazine for free provided that you are a scientist, researcher or engineer who relies on scientific computing in your daily work.

Please [click here](#) to register as a new reader, to update your details or to renew your free subscription.