Towards Photo-realistic 3D Image Capture

Mr Stuart Winsborough
Wicks and Wilson Limited
Morse Road
Basingstoke, Hants
England

Tel: +44 1256 842211
Fax: +44 1256 840997
Email: triform@wwl.co.uk

Contents

Towards Photo-realistic 3D Image Capture

Following last year's presentation on the design and manufacture of 3D scanning systems, this paper concentrates on the capture of colour images using the TriForm 3D BodyScanner system. Current advances in system design that enable more realistic colour models to be produced are presented together with the future technical challenges related to meeting user's expectations of photo-realism.

Issues critical to the capture of colour information in a 3D environment are highlighted, including camera set up, illumination of the subject and resolution of the model. The perceived quality issue – what has to be done to make a colour scan acceptable? – is also discussed.

Different methods of creating and storing colour 3D images, such as point clouds and polygonised models, are presented including a guide to the various file formats currently available.

Finally, applications that make use of such images are introduced. These include applications in the animation and garment industries where visualisation and realism are paramount.

1. Introduction to Wicks and Wilson and the TriForm 3D scanner

Wicks and Wilson was formed in 1973 to manufacture and supply microfilm aperture cards and chemicals to the rapidly growing micrographic industry.

With the advent of the first desktop computers in the early 1980s, the company started to design and manufacture electronic imaging equipment for the micrographic industry. The products included scanners and plotters that used the latest laser and CCD imaging technology.

The company employs 65 people at its Basingstoke headquarters in the UK and is recognised as the market leader in the production of micrographic scanners.

In 1996 a London teaching hospital needing a commercial partner as part of a 3D facial scanning project approached Wicks and Wilson. The core activity of Wicks and Wilson was electronic imaging and it was seen as an ideal opportunity to enter an alternative marketplace.

This project led to the development of TriForm, a method of 3D scanning using a variation of the moiré fringe technique. TriForm systems have since been developed for a number of applications including head scanning and full body scanning with installations at locations in the UK, Europe and the USA.

A dedicated team works on the design and production of the TriForm range of products, with the aim of making it the best human capture system available on the market. In house research and development effort at Wicks and Wilson is focussed on the design of the capture system, while strategic partnerships are being formed with third party companies to service the various vertical markets that can benefit from 3D body scanning.

2. Image Capture

TriForm technology uses a projector and a camera, combined with sophisticated software to capture three-dimensional surfaces. Using this technique, TriForm is able to capture both the shape and colour of a surface.

2.1 Geometry Capture

TriForm uses a structured light technique to capture a surface. This involves projecting a number of patterns of black and white stripes of light onto the surface to be scanned. When viewed through a camera, the stripes that fall on the surface appear to distort or bend. By analysing and comparing a series of such patterns, it is possible to build up an accurate 3D image of the surface.

The system consists of an LCD projector to shine the patterns and a digital CCD camera connected to a frame grabber card in a PC, to record the different patterns as they are rapidly projected onto the surface. It takes approximately 250ms to capture the patterns necessary to calculate the geometry of a single surface.

Each pixel in the CCD is analysed to calculate how far away the surface is from the camera (the z co-ordinate). This data is then combined with the x and y co-ordinates from the camera, which are adjusted to remove optical distortions introduced due to imperfections in the camera and projector lenses. The calculation is repeated for the remaining pixels in the camera to produce a 3D point cloud consisting of x, y and z co-ordinates.

The TriForm BodyScanner captures a total of eight views to give sufficient coverage of the body's surface. The scan sequence takes under ten seconds to capture all eight views that are then analysed to produce the surface geometry of the body.

2.2 Image Capture (Colour)

To understand colour scanning, we must first understand what it is we are trying to capture. What is colour and how is it interpreted by the camera and the human eye?

Light is simply many small packets of energy flying through the air, having been emitted from a source (e.g. a light bulb or the sun). Each packet, or photon, has a certain level of energy and a particular wavelength and frequency. Each of these criteria defines how we see the light, if at all.

White light is made up from equal parts of red, green and blue light. Surfaces appear to have different colours because they reflect or absorb certain parts of the white light. For example a perfectly red surface appears to be red when a pure white light is shone on it because it absorbs all of the blue and green and reflects only the red component. Other colours are produced from a combination of different levels of red, green and blue.

The two extremes are black and white surfaces where the black absorbs all colours and reflects none, whereas the white reflects all colours and absorbs none.

Humans have the ability to interpret colours because at the back of the eye are cells that react to the different red, green and blue components of light. The brain then translates the level of reaction from each of the cells to produce what we understand to be colour.

Digital cameras operate using a similar principal except instead of biological cells, there are CCD cells that produce a electronic signals are processed by a computer before being displayed on a screen.

The human interpretation of colour is nowhere more discerning than in the evaluation of flesh tones and facial tints. It is therefore important for the system to be set up correctly if colour is to be reliably and realistically captured. The camera can be calibrated using a colour chart with specially prepared surfaces that, when pure white light is shone, reflect pure red, pure green and pure blue.

Also important is the way that the surface is illuminated. The light has to be as white as possible and the surface illuminated as evenly as possible. Practical considerations mean that it is not possible to produce perfectly white light, and it is also not possible to illuminate the surface evenly from all directions.

The TriForm BodyScanner uses multiple daylight-matched fluorescent tube lights to produce a diffused illumination across the surface. The lights are built into the scanner in such a way as to illuminate the surface from many directions around the body. The camera then captures the colour data just before the lights are extinguished and the surface geometry is captured, as described above.

2.3 Combining geometry and colour

The geometry is calculated for every pixel in the camera so it is a relatively simple process to match the colour data from the corresponding pixel in the colour map. A point cloud is produced with millions of points, each containing x, y and z geometric data and colour information. Since the same camera is used to capture colour and geometry, perfect registration is guaranteed.

3. Displaying 3D colour images

There are a number of methods for displaying 3D data. The most appropriate method depends entirely on what the images are to be used for and, more importantly, how quickly the data needs to be displayed.

3.1 Point cloud data

The default output from the TriForm system is colour point cloud data. The advantage of point cloud data is that it allows an exact, unmodified record of the original scan to be stored. It is possible to return to a point cloud dataset at a later date if the data needs to be processed in a different way.

However, displaying point cloud data is a processor intensive task, as a body scan may contain 1.5 million points. Even with a high specification PC and an advanced 3D graphics card, it can take about 10 seconds to load the point cloud data. If you do not have a high spec PC, or if you are trying to load the image across a network with limited bandwidth it takes even longer and becomes unacceptable to the normal user.

Point cloud data should be stored as a way of archiving original, unmodified datasets that may require reprocessing in the future, or data that is only accessed on an occasional basis. They should not be used for everyday viewing applications.

3.2.1 Polygonised data - Geometry

Polygonised data differs from point cloud data in that it describes a solid surface, as opposed to a cloud of points. The terms polygonised data refers to any surface made up from polygonal shapes with three, four, five or more sides. The polygonal output from the TriForm system is in the form of triangles.

The easiest way to create polygonised data from a point cloud is to simply ‘join the dots’. If you consider a series of points in the point cloud, joining the dots so that the original points form the vertices of a series of triangles will produce a full resolution polygonised surface.

The production of a full resolution polygonal model does not affect the accuracy of the surface, as all of the original data points are retained. However, the size of the files can be very large. For example a full resolution VRML file can be up to 80Mb in size. Such sizes are totally unmanageable for most applications, except those where surface detail is critical.

It is possible to reduce the size of a polygonised file by using fewer triangles to describe the geometry of the surface. Intelligent point reduction reduces the number of datapoints in the original cloud by using fewer points to describe the flatter surface areas, but keeping a higher density of points in areas of high curvature. This has the effect of reducing the file size without compromising the accuracy of the model too much, making the model suitable for transmission over the internet.

3.2.2 Polygonised data - Colour

Displaying colour information for full resolution polygonised data is a straightforward process. The triangle that is created simply takes the colour of one of its vertices and applies that colour to the whole triangle. The size of the triangles is so small that when viewed on the screen, the model will appear to have high-resolution colour.

However, after reduction the colour mapping becomes more complex. Vertex colour, as described above, can be used, but this will lead to the model having a lower resolution colour and a lower quality appearance that detracts from photo-realsim.

For applications where visual quality is important then a technique known as UV projection can be employed to map multiple colours onto a single triangle. Using this method it is possible to reduce the number of triangles quite drastically, yet retain the fine colour detail at the original capture resolution.

4. File formats

Not every 3D file format has the ability to store colour information. For example, STL files that are common in the reverse engineering sector can only store polygonised surface data. There is no facility for saving colour because that particular market does not require it.

Colour images are most commonly used for visualisation applications, and are becoming more popular with the advent of more advanced internet sites.

VRML has been around for a number of years and is a convenient and relatively universal way of storing the3D colour data. VRML viewers are freely available as plug-ins for web browsers. The disadvantage with VRML is that the entire image has to be downloaded into the remote PC before it is displayed. This can mean a wait of several minutes for internet users with lower spec machines trying to view large models.

To overcome this problem a number of software companies have invented 3D streaming technology. This allows a simplified version of the 3D image to be displayed almost instantly on the remote PC, whilst the rest of the data is ‘streamed’ to improve surface and image quality over time as bandwidth allows.

5. Applications

Not all applications require colour data. For example the health and fitness application, which is described later in the conference, requires only the surface geometry to monitor the change in shape of gym members. Here colour information is irrelevant as it is the shape of the person that is being studied and analysed.

However, visualisation applications do require colour data if they are to be accepted as realistic. Whenever a human body needs to be visualised, the realism of the colour information is critical.

Some of the projects that are currently using TriForm include visualising clothing over the internet and facilitate virtual try-on. For this application it is important to capture the colour of the cloth as well as that of the person.

The animation industry is all about recreating the real world in a virtual environment. Here TriForm allows not only the surface to be captured, but also the colour of skin and clothing. This can save many days work for a CGI artist who would otherwise have to create the material colour from nothing.

6. Future developments

The public has very high expectation levels that are set by exposure to television, computer games and digital photography. If 3D internet applications are to be accepted by the public then the images have to be as real as possible in terms of both geometry and colour. For virtual try-on, the person on the screen has to look like the subject otherwise the application will not be adopted and will eventually fail.

Wicks and Wilson are committed to the quest for photo-realism in 3D imagery. Recent advances in lighting positions and techniques improvements in surface quality and integrity are bringing us even closer to this goal.

One day using a 3D colour image of yourself will be as commonplace as using a digital camera is today. Trying on clothes over the internet or appearing as yourself in a computer game will be the norm. Improving the quality of the images to meet the public’s expectation levels is a key factor in making this possible.

For further information, please visit the Wicks and Wilson web-site at:

www.wwl.co.uk