CHAPTER 6

MISCELLANEOUS ISSUES

Executive summary

This chapter collects together some material on a number of miscellaneous issues such as use of cameras underwater and some practical tips on the use of CCD cameras.

6.1 Underwater camera calibration

There are two basic modes of underwater photogrammetry. These involve either cameras designed specifically for underwater work or, cameras usually used in air which are enclosed in a waterproof housing. Those designed for direct underwater work, such as the popular 35 mm Nikonos range, have had their lenses manufactured on the understanding that they would be in direct contact with the water. That is, the incoming light rays are assumed to be coming from a medium with a refractive index close to 1.34. Acknowledging this factor means that the underwater lenses can produce images with only about as much radial and decentering distortion as an ordinary lens would produce for an image of an object in air. If those underwater lenses are used in air, large distortions occur.

The other option for an underwater photographer is to use a 'normal' camera and lens and encase it in an underwater housing. These housing are of two types: those which have been specially designed for a type of camera/lens and have a spherical glass dome on the front of the housing and those which are more 'universal' in application and can be considered 'box-like' with a plane glass port. Those with spherical domes tend to dramatically reduce the impact of radial lens distortion. As an example of those with a plane glass port, the radial distortion an ordinary 35mm camera placed into such a housing may increase from an in-air value of, say 100 µm at the edges of the 36 x 24 mm format to an in-water value of approximately 1500 µm!! If a proper calibration is carried out, even distortion errors of this magnitude can be corrected with confidence. Underwater housing with spherical ports are often quite expensive, so those with a parallel glass port are often used.

It can be realised, even from this short introduction, that it is radial distortion which causes 'problems' in multi-media situations. The plane glass port simply acts like another lens element, and since it may be significantly thick (10 mm is not uncommon) to withstand the pressures involved with depth, and physically separated from the lens, the effect of radial distortion is highlighted. Should the camera's axis not be placed perpendicular to the glass port, then significant decentering distortion may also become apparent.

A note on underwater stereo pairs

The advantages of obtaining a stereo-pair of photographs to enable the exposition of some object details are well-known and this technique is commonly used for in-air situations. A note of caution must be given for those contemplating the same maneouvre with underwater photographs. In order to re-create an underwater stereo-scene, unless an analytical or digital stereoplotter is used which has appropriate software to compensate for the large lens distortions and the refractive effects of water, there will be uncomfortable amounts of 'y-parallax' and a problem in trying to define heights in the stereo-model.

Consider the moment of stereo-photography: the light rays coming from an object will travel straight to the lens (assumed to be an ordinary in-air lens for this discussion), but then will be refracted to a position on the image which is not in an equivalent straight line passing through the centre of the lens, but is radially distorted away from the centre of the image. Upon viewing the stereopair under a stereoscope, the rays will be re-constructed from that radially distorted position to meet at a 3-D location considerably closer to the camera lens than is true. The viewing is taking place in air and there are no compensating effects of refraction.

More disturbing than the height problem can be the fact that most pairs of light rays will no longer intersect in the re-created model. Only in the special situations where the object is equidistant from each lens will corresponding rays intersect in the model. In all other cases the observer will have to accept a 'compromise' when trying to view the underwater scene in 3-D. As mentioned above, analytical and digital stereoplotters can be programmed to compensate for this effect.

The scenario described above is perhaps the simplest of the multi-media situations in which photogrammetrists may find themselves. Consider an experiment behind a glass panel where the objects are immersed in a fluid/liquid which may or may not be water. If it is necessary to stand back a distance from the glass shield in order to gain a large enough field of view, then the distances object-to-glass and glass-to-camera are both important in correcting for distortion. A simple example would be studying movement patterns of a trawling net being towed through a test tank which was fitted with glass sides. Other examples would include aerospace components subjected to intense cold/heat conditions and photographed through the observing panel of a thermal chamber. To allow viewing through a diversity of photographic angles, it is usual to place the object to be examined on a rotatable turntable in these chambers.

Appropriate Calibration Techniques

The methods of plumbline calibration and self-calibration have been used to great effect in multi-media situations. A plumbline calibration can be used to isolate and obtain excellent values for the parameters of radial and decentering distortion before these values are held 'fixed' in a self-calibrating bundle adjustment which should provide values for the principal distance and the offsets of the principal point.

The refractive index of water changes only by a very slight amount (less than 1%) through a range of temperature, pressure and salinity conditions. In fact, unless the depth is greater than 1000 metres (say) it can be considered as a constant. This means that plumbline and self-calibration tests can be undertaken in a swimming pool for experiments which may be hundreds of metres under the surface of the ocean. In the case of underwater photogrammetry, the camera to glass port distance will remain fixed, so the one set of camera parameters will suffice for all exposures.

In the case of the thermal chamber photography where the object will be rotated on a turntable, the camera can be secured on a tripod at a convenient distance from the viewing panel, so that this geometry remains constant throughout the photography. This is a desirable condition, all images will be affected by the same amounts of radial and decentering distortion.

Some notes of caution

When cameras are used underwater, it has been the author's experience that something usually 'goes wrong'. Water is an unforgiving environment for electronics. Usually problems arise with flash units, as these are extremely susceptible to moisture problems. Technicians will religiously check and grease the 'O-rings' on the cameras (if of the underwater immersible type) and on the camera housings, but only pay cursory attention to the flash units.

Daylight does not penetrate ocean water effectively to depths greater than 10 metres, and without a flash at such depths, the 'red' component of all objects will literally disappear from imagery. Consequently, don't make targets of any red colour, and preferably use black and white. Of course retro-reflective targets start to lose their effectiveness after a depth of a few millimetres!

The earlier discussion in this section about radial distortion effects due to the refractive index of water did not make specific mention about the principal distance, but of course it is directly affected by the water. Just as gold-fish in a bowl appear to be 34% larger than they are, so it is when designing an underwater survey that care must be made when deciding which lens to choose. A 28 mm lens will act like a 38 mm lens underwater. In general, try to use a wide angled lens because not only will it act like a lens of longer focal length but you will then have the capability of getting closer to the object in order to image it. The clarity of water is not something the photogrammetrist can choose and noting the rapid attenuation of light underwater, it is important to be as close to the object as possible.

A final factor to consider is the effect that a flash can have in an underwater situation. In air, it is usual to place the flash as close as possible to the lens for photogrammetric work. Even ring strobes are placed around lenses. Underwater, this is usually a disaster as the fine elements of particulate matter which seem to be in even the clearest water will cause much of the flash to rebound to the lens. The further (within reason) that the flash is from the lens, the better in underwater situations. The majority of the rebounding light from the particles in the water will aim straight back to the flash, so if the camera is offset, a much better image will be recorded.

6.2 Useful tips

The following sections include some tips which practitioners may find useful when using electronic cameras for photogrammetric purposes.

A question which is often asked concerns the size of CCD pixels. Since the accuracy of a final result one can achieve with an electronic camera is (almost) directly linked to the number of pixels on the array, why aren't the arrays made larger? The answer to this concerns the manufacturing process for CCD arrays. It is simply incredibly difficult (and hence prohibitively expensive) to produce arrays with lots and lots of pixels. For examples: a 500 x 700 pixel array is of the type mass-produced for video camcorders and such a device sells for under 100 pounds. A 1000 x 1000 array (so-called Megapixel) presently will cost 3,000 pounds while a 2,000 x 3,000 pixel sensor will need a bank loan for 15,000 pounds (prices as at mid-1998). The largest sensor array available commercially at present is 4,000 x 7,000 and if the manufacturer will sell a single chip to anyone is doubtful as the entire production goes to large specialist medical imaging companies, but it costs over 50,000 pounds! The larger the sensor, the larger the number of defective pixels per chip, so much so that perhaps 99% of a production run is so badly degraded as to be re-cycled!

Another common query is why aren't the pixels made smaller in size? Everyone knows that the accuracy of centroiding targets is about one thirtieth of a pixel (say 0.3 micrometres), so the smaller the pixels, the more accurate the result! There are physical limitations to the amount by which pixels can be "shrunk". The quality of the lens and the wavelength of light define the smallest spot size of light which can be distinguished on the image. The amount of light (number of photons) which can be collected in a short time frame is another consideration. It simply is not worthwhile to try and produce pixels smaller than a physical limit of three micrometres, although the trend in the past decade has seen pixels reduce from approximately 15μm to 8μm.

Digital cameras overcome many of the effects displayed by analog video cameras, but a mention of some of these is warranted for completeness. The analog signal produced by a video camera includes horizontal and vertical synchronisation information which is used during sampling of the image for analog-to-digital conversion as well as information about the image itself. The geometric quality of an analog video image depends on the synchronisation capacity of the frame-grabber. If the synchronisation information, or pixel-clock, is driven from the camera, better results will be obtained. A phenomena termed line-jitter can exist when the frame-grabber samples the camera's output signal at a rate which is slightly different to that of the camera. Other problems which can occur with an analog video camera/frame-grabber system include phase patterns and the introduction of noise into the image due to electronic interference in the transmission of the signal.

Many first-time users of frame-grabber systems are totally unaware of some of the geometric changes which can occur between camera capture and image display. It is not uncommon for a camera with a quoted pixel array of (for example) size 568 x 748 pixels (each purported to be 9μm square) to omit a surrounding annulus of 4 pixels and only transmit an analog signal of 560 rows each composed of 740 pixels. The frame grabber in the computer may resample this signal to (say) 550 x 700 pixels. Did the frame-grabber leave out the first 10 rows, the last 10 rows, some other combination or merely do some mathematical resampling? What about the number of columns? Most likely, each row has been re-sampled so most of the data will be across the image, but what does this do the principal point? Where is it now? When viewed on a PC-screen, the pixels may well be converted to an elongated shape, perhaps in the ratio 9 x 6. Of course, most users of photogrammetry just want answers for 3-D coordinates and fortunately most software can accommodate these changes, but before using new equipment, some inquiries such as these should be made.

Changes in temperature can have significant effects on image geometry. These are referred to as warm-up effects. Most researchers state the equipment should be switched on for 90 minutes before use. There are two likely causes of warm-up effects: firstly, the expansion of the sensor due to temperature increase and secondly, the variation of the internal clock frequency.

Image compression, such as the well-known JPEG, is something for which photogrammetrists must be wary. JPEG (and other so-called ‘lossy’ compression techniques), replace the actual grey-scale values of the pixels with mathematical functions in a patch-wise pattern across the image. Depending on the amount of compression, the number of mathematical coefficients which are stored will vary. Clearly, upon restoration of a patch of an image from a mathematical model which has sampled that patch, high frequency effects will be lost and photogrammetrically important information such as edges may suffer sub-pixel movement. Since targets can be located to one or two tenths of a pixel, JPEG and other compression techniques must be avoided. Unfortunately it is sometimes very difficult, if not impossible, to discover what some proprietary digital cameras do to their images immediately after capture and before downloading takes place.

Illumination is an important component of any imaging system. Diffuse lighting usually is preferred to direct illumination. Sources which emit large amounts of infra-red radiation should not be used, as CCD arrays are sensitive to them. Of course, discussion of illumination is only relevant to the type of targets being used. Assuming retro-reflective material is to be used (other possibilities include laser, naturally reflective or projected targets), it has been found that masking the retro-reflective tape to produce the target can cause the target position to shift systematically when viewed from varying directions.

The centroiding technique to be used which “best” finds the true centroid of a circular target has been the subject of considerable investigation. Some options include ellipse fitting, least squares template matching, using only a binary centroiding algorithm where only pixels with grey-levels above a certain threshold were considered, and centroiding based on a weighted centred of gravity method (weighted by either the grey value or the grey value squared of the pixels). The algorithm chosen by most researchers is the weighted centre of gravity method where the weight is only the grey value (squaring the grey value tended to exaggerate errors due to noise).

The size of target is also a consideration as researchers have found the size of target to influence the precision of its location. Precision improves as target diameter increases to six pixels. Due to some cameras being further from the object than others, and all targets not being optimally located for any one image, a range of 5 to 10 pixels is suggested when designing targets.

Test-ranges still play a useful role in the calibration of camera/lens combinations. It has already been expressed that three-dimensional test ranges are expensive to construct and maintain. Of course, for several tasks they are ideal, but much good work can be done with a simple two-dimensional test area. For example, a set of retro-reflective targets placed on a flat wall in a laboratory can be most useful. As long as images are captured from a wide variety of convergent directions and the camera is rolled through 90 degrees in order to 'projectively uncouple' some of the parameters in the calibration model, then good results can be obtained. At least 30 targets are suggested and about 8 or more images should provide sufficient information for a reasonable result. The method of self-calibration means that the targets do not need to be coordinated prior to photography. The addition of two scale bars at right angles to one another certainly strengthens this arrangement.