

ACCIDENT RECONSTRUCTION VIA DIGITAL CLOSE-RANGE PHOTOGRAMMETRY

Clive Fraser, Professor
Department of Geomatics
University of Melbourne
Victoria 3010, Australia
c.fraser@unimelb.edu.au

ABSTRACT

With the ever wider adoption of spatial information analysis, there has been an increasing awareness of the importance of 3D data acquisition systems to support traffic accident reconstruction and forensic analysis. In an application such as accident reconstruction, GIS and CAD systems have limited utility without the underlying data. Digital close-range photogrammetry displays many attributes which make it a well suited technology for the provision of the necessary 3D measurement data. The application of photogrammetry to accident reconstruction is not without difficulties, however. These include accommodation of the generally very poor, near-planar network geometries encountered and the need for maximum ease of use, from which follows the requirement for highly automated processing and fully automatic camera calibration. This paper concentrates upon developments undertaken to enhance the applicability of close-range photogrammetry and consumer-grade digital cameras to accident reconstruction and forensic measurement. The developments comprise a new approach for robust on-line image orientation and a method for automatic camera calibration which employs colour coded targets. These are highlighted through reference to the *iWitness* system for close-range photogrammetry.

INTRODUCTION

The 3D mapping of traffic accident scenes, which is commonly termed Accident Reconstruction (AR), is employed for a number of reasons. These range from the support of technical investigations, such as vehicle speed determination and analysis of vehicle collision event dynamics, to the provision of evidence in court hearings. Whereas a vehicle manufacturer or traffic engineer might need a detailed 3D reconstruction, a local police force may only require simple 2D documentation in recognition of the fact that if the accident does not result in subsequent legal proceedings, then the AR data will likely never be used. Irrespective of the comprehensiveness of the required 'mapping', the essential first step in the process is to accurately characterize the dimensions of the accident scene. The final outcome of the AR process is typically a CAD drawing in the first instance, which may be further developed into a 3D model or even a computer animation.

For example, from the accident scene shown in Fig. 1, a simple 2D 'diagramming' is first performed, as shown in Fig. 2, after which 3D views, as exemplified by Fig. 3, may or may not be required. It is noteworthy that limited sophistication is required in AR; the aim is to obtain a dimensionally accurate 3D representation of the scene, rather than a comprehensively rendered or textured model. The CAD software that produced Figs. 2 and 3, namely Crash Zone (CAD Zone, 2006) costs less than \$1000. Both of the drawings shown could be adequately accomplished with 2D surveying, at its simplest represented by the measurement of distances along and offset from a 'baseline' (e.g. road edge or centreline), as was traditionally done. However, with the enhanced scrutiny of any evidence in a court, and the need for the AR data collection process to be as least disruptive to traffic as possible, the requirement has arisen for more complete and accurate data to be recorded in the shortest time possible. Over recent years, total stations, laser range finders with angle encoders, and even laser scanners have been used. These technologies have resulted in more comprehensive 3D modeling, but not necessarily faster data acquisition at the accident scene. Moreover, they are relatively expensive and complex for local police and traffic agencies.

The number of law enforcement agencies in the US is estimated to exceed 15,000, with agencies ranging from city and county police to state highway patrols. The majority of these have some level of requirement for AR and forensic measurement, and thus there is a demonstrated need for an AR measurement technology which can offer very low-cost, flexible mapping of accidents with an absolute minimum of on-scene recording time. These imperatives have seen attention turn to close-range photogrammetry. Indeed, the low-cost photogrammetric software

suite *iWitness* (Fraser & Hanley, 2004; Photometrix, 2006; Fraser et al., 2005) has been developed primarily for AR and forensic measurement. AR displays some distinctive characteristics and difficulties within a photogrammetric measurement context, and these call for special attention when designing a purpose-built close-range photogrammetric system. This paper will discuss some of the necessary innovations and developments required to render photogrammetry suitable for AR.



Figure 1. Example accident scene to be reconstructed (note that traffic cones serve as object feature points).

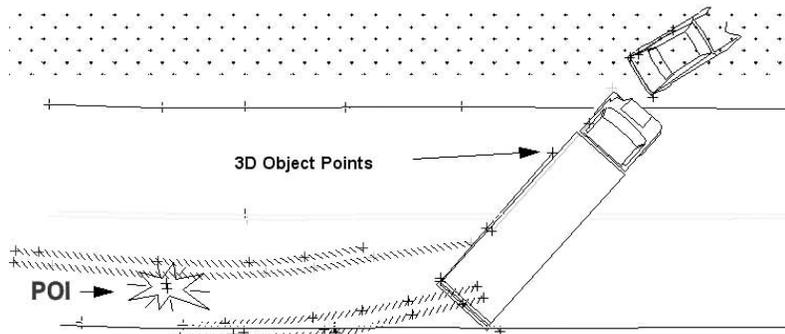


Figure 2. Example CAD drawing of accident of Fig. 1, illustrating object features of interest.

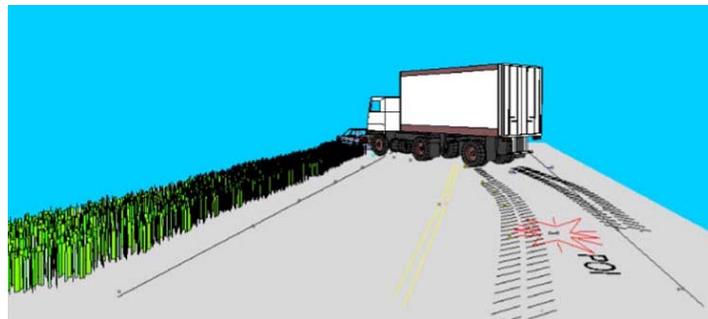


Figure 3. View of 3D model of accident of Fig. 1.

iWitness OVERVIEW

The *iWitness* system for digital close-range photogrammetry was primarily designed for AR and forensic measurement. It is characterised by innovations within its image measurement and photogrammetric orientation processes which have been expressly developed to address some of the difficulties encountered in AR. A short overview of *iWitness* will be presented here, after which two developments to enhance the application of affordable close-range photogrammetry to AR will be detailed. The first of these concerns initial network orientation, which is greatly complicated by the near-planar object point fields encountered in AR. The second is fully automatic camera calibration for the consumer-grade digital cameras that are employed with *iWitness*.

iWitness generates attributed point clouds, with the attributes primarily being lines which are preserved in the export of object coordinate data in DXF format. The system is designed to interface with CAD and modeling packages. The graphical user interface of *iWitness* is illustrated in Fig. 4, which shows the vehicle collision survey from which the 'diagramming' shown in Figs. 2 and 3 was produced. The *iWitness* system has many features over and above the orientation and calibration developments that are discussed here. These include automatic recognition of the camera(s) via information contained within the EXIF header of the JPEG or TIFF images, and fully automatic initiation of all computational functions, i.e. computations are never specifically invoked but occur automatically in the background with every image point 'referencing'.

Also included is a 'Review Mode' whereby it is possible to interactively review all image point observations and to adjust these where appropriate, again with on-line and immediate updating of the photogrammetric bundle adjustment. A quality measure, which provides an effective error detection and correction capability, indicates any subsequent improvement or degradation in the spatial intersection accuracy as the review process is undertaken. Semi-automatic image point measurement of artificial and even some natural targets via centroiding to an accuracy of up to 0.03 pixels is also available.

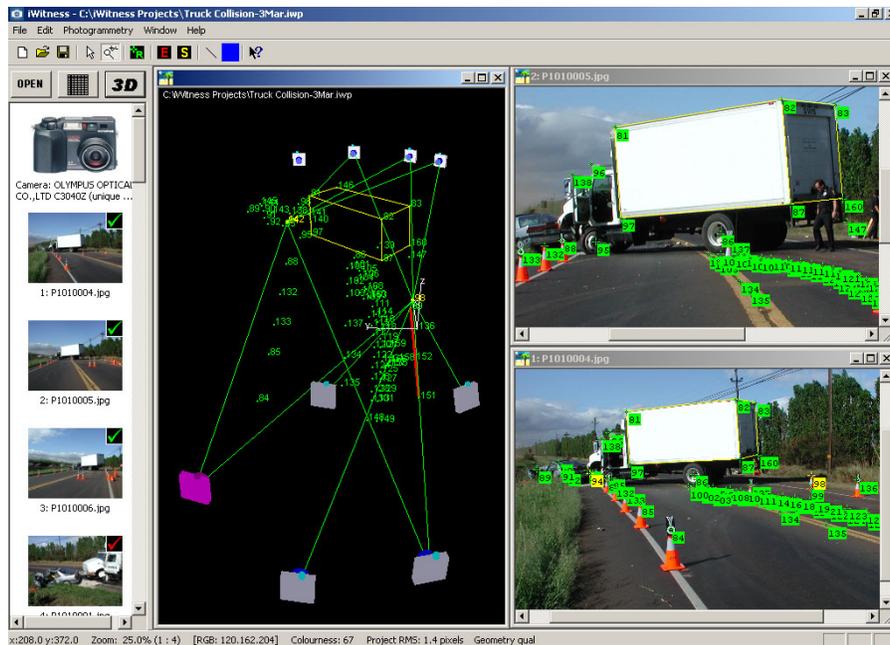


Figure 4. *iWitness* user interface; the CAD diagrams in Figs. 2 and 3 are from this survey.

CAMERA STATION CONFIGURATIONS IN AR

Perhaps the most significant problem that faces someone wishing to apply photogrammetry to AR is that, by and large, all feature points of interest tend to be near planar in their distribution, the majority lying on or near the road surface. A traffic accident scene can be 50-200m or more in length, but often displays a vertical range of interest of only a few metres or less. Long and thin near-planar object point arrays constitute a quite undesirable

geometric network configuration for close-range photogrammetry. Exacerbating this problem is the fact that the camera stations also lie close to the plane containing the target points. This is well illustrated in Fig. 5, which is both a real and generally representative AR network. The near planarity of camera stations and object points is clearly apparent, as is the fact that the network is long and narrow. Photogrammetric surveys for AR need to be conducted without any requirement for control points, and it can be seen that the camera station geometry in Fig. 5 is very unfavourable from the point of view of both building an initial relative orientation (irrespective of the chosen image pair) and subsequently carrying out a multi-image bundle adjustment.

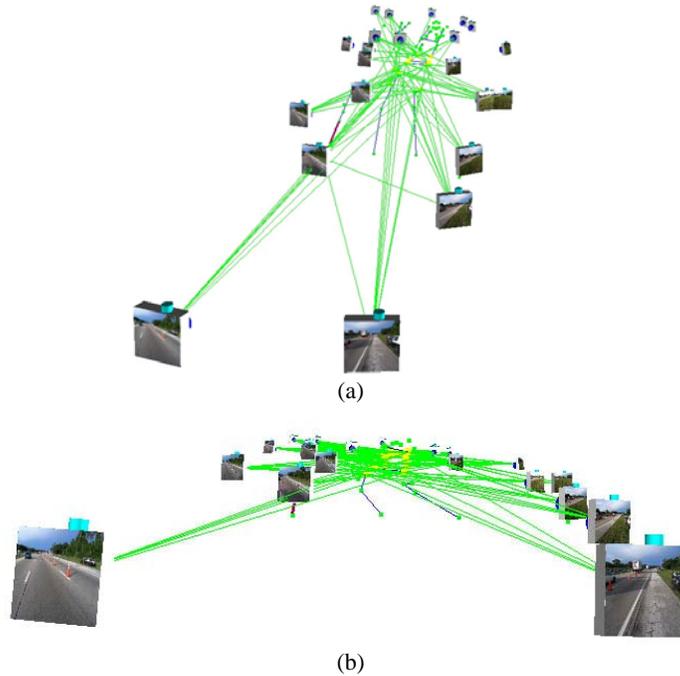


Figure 5. Two views of an unfavourable, yet typical network geometry encountered in AR.

Although AR networks do not usually employ any ground control, though they are scaled, special targets can be used to alleviate to some degree the problem of poor geometry. For example, it is common to employ ‘evidence markers’ (DCS, 2006), which are back-to-back targets, as illustrated in Fig. 6. These face horizontally and can be semi-automatically measured in *iWitness* via an operated-assisted centroiding function. While evidence markers facilitate accurate conjugate point referencing from opposite directions, they do nothing to enhance the otherwise weak network geometry. The near-planar point distribution can be overcome by, for example, feature points on the vehicles involved, street signs, traffic cones and even tripods. However, the fact remains that from a photogrammetric perspective the most challenging part of AR applications is network orientation. To conquer this problem, *iWitness* incorporates some innovative orientation procedures, especially for relative orientation.



Figure 6. Evidence markers placed on features of interest.

ON-LINE PHOTOGRAMMETRIC ORIENTATION

As highlighted above, photogrammetric network geometry in AR can be complex; indeed far more so from a sensor orientation standpoint than the stereo geometry of topographic photogrammetry or the binocular stereo or wide baseline geometries encountered in computer vision. Coupled with the often highly-convergent and multi-magnification camera station arrangements are object point geometries which are generally unsuited to relative orientation (RO) and spatial resection.

There are effectively two basic mathematical models for sensor orientation: the coplanarity equation for RO, and the collinearity equations for spatial resection, intersection and multi-image bundle adjustment (exterior orientation) with or without camera self-calibration. Both constitute parametric models which, in their linearized form, are solved via an iterative least-squares adjustment of initial values for the parameters. In the on-line computational process of *iWitness*, where the least-squares bundle adjustment is updated as each new observation is made, it is imperative that the initial values of exterior orientation parameters are determined with sufficient accuracy to ensure solution convergence.

Two approaches have traditionally been adopted for the determination of preliminary exterior orientation. The first of these involves the use of object space 'control points' with known or assigned XYZ coordinate values. These points, which need to number four or more in at least two images in the network, then facilitate closed-form spatial resection. Spatial intersection can follow to establish the object coordinates of further image points, which in turn can support resection of further images and spatial intersection of additional points, and so on. Nowadays, the use of exterior orientation (EO) devices is popular in industrial vision metrology systems as a practical means of providing the necessary 3D control points for automated initial exterior orientation (Fraser, 1997; Ganci & Handley, 1998).

A second approach, which has not been widely adopted, is initial RO. The attractiveness of RO is simply that it requires no object space coordinate data. Moreover, it is well suited to image measurement scenarios where conjugate points are 'referenced' between two images, point by point, as within a stereoscopic model for example. It is well known that for a given image pair a minimum of five referenced points is required to solve for the unknown parameters in a dependent RO via the coplanarity model.

It is also well established that for convergent imaging geometry good initial parameter approximations are required to ensure convergence of the RO solution. With the addition of the third and subsequent images, resection would follow. Here too, good starting values are necessary, though unlike the situation with RO, there are well recognised closed-form and two-stage solutions for resection. The most pressing problem encountered in developing a robust, reliable solution for RO in *iWitness* was finding a method for generating initial values for the five RO parameters of rotation (3) and relative translation (2). Experience with the least-squares solution of the coplanarity equation suggests that it is very stable when representative initial parameter values are available, even in situations of very poor geometry.

Alternative models for image orientation have been adopted by the computer vision community. The most popular of these is the Essential Matrix formulation for solving in a linear manner the position and orientation of one camera with respect to another, which was introduced to computer vision by Longuet-Higgins (1981). The essential matrix formulation implicitly assumes 'calibrated' cameras, or in photogrammetric terms, known interior orientation. An 'uncalibrated' version of the essential matrix is the Fundamental Matrix (Hartley & Zissermann, 2000). Upon a review of the literature, the impression is certainly gained that these approaches have considerable promise as a means to solve RO. This is notwithstanding concerns that linear solutions for the essential and fundamental matrices are prone to ill-conditioning and the generation of both erroneous solutions and matrices which are not always decomposable. Regrettably, while there are many publications dealing with theoretical and algorithmic aspects of the essential matrix approach, there are not too many that give a comprehensive experimental analysis of the method, especially in cases of poor geometry.

Thus, an evaluation of the essential matrix approach for the estimation of initial RO parameters in *iWitness* was undertaken. The endeavour, however, was unsuccessful from the standpoint of producing a robust, scene independent RO solution that would be amenable to later refinement via the rigorous coplanarity model. The prospect of success with near-planar objects could immediately be discounted, since this is a known failure case – but a geometry that is unfortunately prevalent in AR. In order to enhance the prospects of success, a normalization process, RANSAC approach and singular value decomposition (and possibly two) can also be applied. Regardless of what approaches were implemented, however, it was found that the method was unreliable and unstable for an application demanding at least a 95% success rate. An alternative was therefore sought for *iWitness*. A quote from Horn (1990) is noteworthy at this point: "Overall, it seems that the two-step approach to relative orientation, where one first determines an essential matrix, is the source of both limitations and confusion".

The search for a robust procedure for RO in *iWitness* ended with the adoption of a Monte Carlo type strategy whereby a very large number of possible relative orientations are assessed for the available image point pairs. The refined solution in each case is obtained via the coplanarity model using combinations of plausible initial values (there could be hundreds of these). From the number of qualifying solutions obtained for the first five point pairs, the most plausible are retained. But, no RO results are reported to the user at this time, as there may be quite a number in cases of weak geometry, compounded by noisy data, therefore leading to the likelihood of ambiguous solutions. This process takes only a fraction of a second. Then, as point pairs are successively observed, the computation is repeated, with the aim being to isolate the most probable solution from the ever fewer qualifying candidates. Once there is a sufficient degree of certainty as to the correct solution, the orientation computation swings from a coplanarity to a collinearity model, namely to a bundle adjustment. In cases of reasonable network geometry and camera calibration, a successful RO is typically reported to the operator after seven point pairs are 'referenced'. For weaker geometry and/or very poor calibration, the number of required point pairs may rise to eight or nine and on occasion to more than 10.

A similar approach to checking plausible orientation solutions on line is employed when new images are added to an already oriented network. This time, spatial resection computations are performed via a closed-form algorithm. Generally, the criteria for a correct solution are met after five or six point pairs are referenced, though in favourable cases only four points are required. Once resection is successful, the image is added to the network and on-line bundle adjustment is used to integrate subsequent image point observations. This unique approach to on-line exterior orientation is a very powerful and popular feature of *iWitness* since it is robust, very well suited to blunder detection, and occurs instantly and automatically.

AUTOMATIC CAMERA CALIBRATION

It is well established that for camera self-calibration a multi-image, convergent camera station geometry, which incorporates orthogonal camera roll angles, is required along with an object point array which yields well distributed points throughout the format of the images. Initial starting values for the camera calibration parameters are also required, though with the exception of the focal length, these initial values may be taken as zero. The accurate modelling of lens distortion is assisted by having well distributed image points throughout the image format.

By taking advantage of the robust exterior orientation process described in the previous section, a network for self-calibration can be provided through the provision of image point correspondences alone, i.e. from the (x, y) image coordinates for all matching points. The approach taken in *iWitness* to ensure fast and accurate matching of image point features is based on colour coded targets (Cronk et al., 2006). Traditionally, coded targets employed in close-range photogrammetry have been geometric arrangements of white dots or shapes on a black background (Fraser, 1997). These targets require optimal exposure to ensure a near binary image is obtained and although such a requirement may be practical for the controlled environments of industrial photogrammetry, it does not suit the conditions encountered in AR. Nor does it take advantage of one of the most prominent characteristics of today's digital cameras, namely that they produce colour (RGB) imagery.

Fig. 7 shows the colour codes designed to facilitate fully automatic camera calibration in *iWitness*. The 5-dot geometric pattern of points is the same for each target, but the colour arrangement of red and green dots varies so as to yield 32 (2^5) distinct codes. Once the code dots are detected, a colour transformation process is used to isolate the red/green arrangement and identify the code. The adoption of colour codes has afforded a more flexible approach to automatic self-calibration.

It is usually convenient to simply place the codes on the floor, with one or more being out of plane, as indicated in Fig. 7. Non-planarity of codes is not essential, but generally aids in both the initial network orientation and in reducing projective coupling between the interior and exterior orientation parameters. The precision of the recovered calibration is thus enhanced. As previously mentioned, an initial value for focal length is required, however this is not really the case for the operational system. The procedure again follows a trial and error scenario where multiple principal distance values are tested as the network is being formed and the most plausible value is taken as the initial estimate within the final self-calibrating bundle adjustment.

A typical network for automatic calibration based on colour codes is indicated in Fig. 7. The codes are purposefully chosen to be relatively large, not only to aid in recognition, but also to constitute a sub-group of points. Rather than being treated as a single point, each code forms a bundle of five rays, as seen in the figure. A broader distribution of image point locations is therefore achieved, which adds strength to the photogrammetric network.

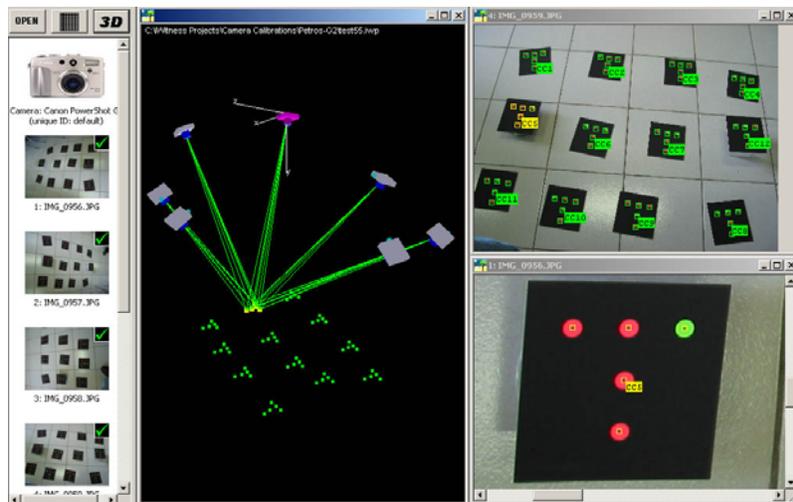


Figure 7. Automatic camera calibration in *iWitness*; note array of 12 colour coded targets.

CONCLUDING REMARKS

This paper has discussed the application of digital close-range photogrammetry to AR, with the principal topics being the two innovations of robust on-line image orientation and automatic camera calibration. These developments, which are incorporated in the *iWitness* software system, enhance the utility, robustness and flexibility of digital close-range photogrammetric applications employing off-the-shelf cameras. Although the development of the new exterior orientation process and automatic calibration via colour coded targets was driven by the needs of the AR and forensic measurement sector, these innovations are equally applicable to a wide range of close-range, image-based 3D measurement tasks. The combination of *iWitness* and an off-the-shelf digital camera of greater than 3-megapixel resolution affords prospective users of close-range photogrammetry the ability to undertake measurement tasks requiring accuracies of anywhere from 1:1000 to better than 1:30,000 of the size of the object, for as little as \$2000.

REFERENCES

- CAD Zone (2006) Company web site: www.cadzone.com (site accessed Feb. 2, 2006).
- Cronk, S., Fraser, C.S. and Hanley, H.B. (2006) Automatic calibration of colour digital cameras. *Photogramm. Record*, 19 pages (under review).
- DCS (2006) DeChant Consulting Services – DCS, Inc.: www.photomeasure.com (site accessed Feb. 2, 2006).
- Fraser, C.S. (1997) Innovations in automation for vision metrology systems, *Photogramm. Record*, 15(90), 901-911.
- Fraser, C.S and Hanley, H.B. (2004). Developments in close-range photogrammetry for 3D modelling: the *iWitness* example. *International Workshop: Processing and Visualization using High-Resolution Imagery*, Pitsanulok, Thailand, 18-20 November, 4 pages (on CD-ROM).
- Fraser, C.S., Hanley, H.B. and Cronk, S. (2005) Close-range photogrammetry for accident reconstruction. *Optical 3D Measurements VII*, (Gruen, A./Kahmen, H, Eds), Vol. II, 115-123.
- Ganci, G. and Hanley, H. (1998) Automation in videogrammetry, *International Archives of Photogrammetry and Remote Sensing*, 32(5), 53-58.
- Hartley, R.I. and Zissermann, A. (2000) *Multiple View Geometry in Computer Vision*, Cambridge Press.
- Horn, B.K.P. (1990) Recovering baseline and orientation from essential matrix, <http://ocw.mit.edu/OcwWeb/Electrical-Engineering-and-Computer-Science/6-801Fall-2004/Readings/>, 10 pages.
- Longuet-Higgins, H.C. (1981) A computer algorithm for reconstructing a scene from two projections, *Nature*, 293, 133-135, 1981.
- Photometrix (2006). Website for *iWitness* system: www.photometrix.com.au (site accessed 2 Feb. 2006).