CLOSE-RANGE PHOTOGRAMMETRY FOR ACCIDENT RECONSTRUCTION

Clive Fraser, Harry Hanley and Simon Cronk
Department of Geomatics
University of Melbourne
Victoria 3010 Australia
Email: (c.fraser, hhanley, cronks)@unimelb.edu.au

Abstract: Throughout the last decade forensic scientists, technicians and police have employed a number of 3D measurement tools for crime scene and accident reconstruction. These have ranged from the basic, such as EDM instruments, to the complex, namely terrestrial laser scanners. In the field of traffic accident reconstruction, close-range photogrammetry is now being adopted, primarily because of the greatly reduced on-scene time, which leads to shorter periods of traffic disruption. The fact that a permanent visual record is also obtained, from which 3D measurements can be made at any time, is a further notable benefit. However, for successful application of close-range photogrammetric techniques in accident reconstruction a few important issues must first be dealt with. These include accommodation of the generally very poor, near-planar network geometry 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 reports upon two innovative developments undertaken to enhance the applicability of close-range photogrammetry and consumer-grade digital cameras to accident reconstruction. The developments comprise a new approach to robust on-line image orientation and a method for automatic camera calibration which employs colour coded targets. They are highlighted via the *iWitness* system, which has been developed primarily for accident scene reconstruction and forensic measurement applications.

1. INTRODUCTION

The aim of traffic accident reconstruction (AR) is, as the name implies, to reconstruct motor vehicle collision scenes. Whether the final requirements of the AR process are to assist in calculations (such as vehicle speed), to analyse the dynamics of the collision event(s), to provide evidence in a subsequent court case, or for some other purpose, an essential first step is to accurately characterise the dimensions of the accident scene. The comprehensiveness required can vary depending upon the ultimate use of the 'mapping' data produced. For example, a vehicle manufacturer or traffic engineer might need a detailed 3D reconstruction, while the 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. Unfortunately, it is not always known at the time of the accident whether court proceedings will eventuate. In most jurisdictions, accidents involving fatalities must be surveyed and mapped. The term 'diagramming' is used in the US to describe this documentation process, since the final outcome is typically a CAD drawing in the first instance, which may be further developed into a 3D model and even an animation.

Shown in Figs. 1 and 2 are examples of CAD drawings for two accident scenes. In the context of 3D modelling, both are reasonably simple representations. Also, both could be adequately accomplished with 2D surveying, at its simplest represented by the measurement of distances along and offset from a 'baseline' (eg. 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 comprehensive and accurate data to be recorded in the shortest time possible. More recently, total stations, laser range finders with angle encoders and even laser scanners have been used. In the case of expensive laser scanning technology, however, adoption has been mainly confined to research laboratories and large centralised accident investigation agencies. These technologies have resulted in more comprehensive 3D modelling, but not necessarily faster data acquisition at the accident scene. Moreover they are relatively expensive and complex for local police and traffic agencies.

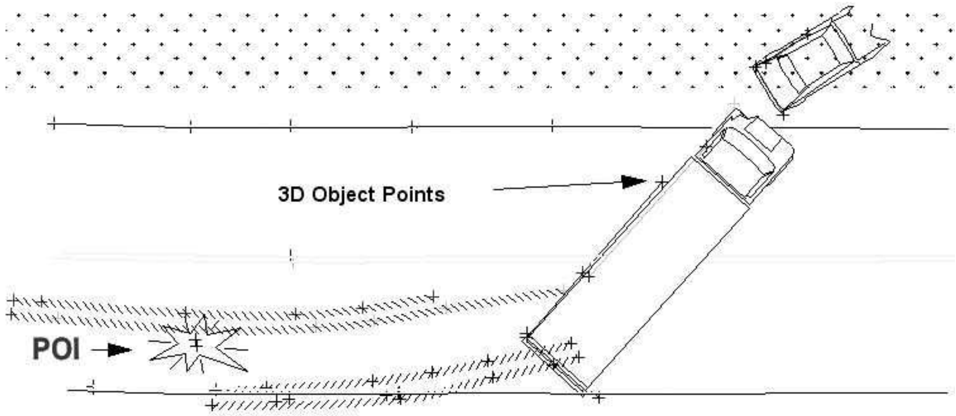

Figure 1: Example CAD drawing for AR illustrating object features of interest (courtesy of DeChant Consulting Sevices – DCS [2]).

In the US alone, there are in excess of 10,000 law enforcement agencies. In relation to AR these range from city and county police to state highway patrols. For the large number of local agencies involved in AR, a technology is needed that 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, a low-cost photogrammetric software suite, called *iWitness* [10], has been designed and developed primarily for AR and forensic measurement. Our purpose in this paper is not so much to extol the virtues of photogrammetry to readers who are quite familiar with the technology, but rather to consider some of the distinctive characteristics of AR that call for special attention when designing a purpose-built close-range photogrammetric system.

2. *iWitness* OVERVIEW

The *iWitness* system is characterised by a new paradigm within its image measurement and photogrammetric, namely automatic on-line computations which are never specifically invoked but occur automatically in the background with every image point 'referencing'. We will present a short overview of *iWitness*, after which we will concentrate on two developments to enhance the application of affordable close-range photogrammetry to AR.

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*.

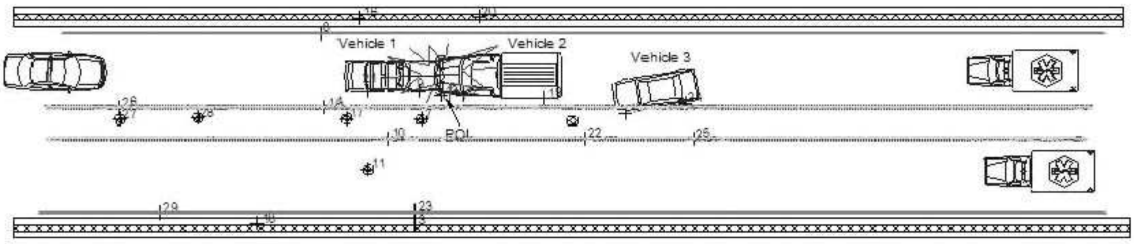

(a) plan view

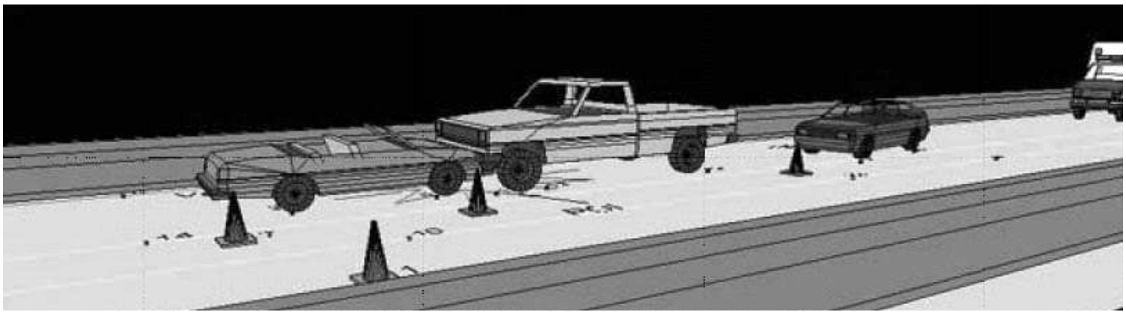

(b) perspective view of model

Figure 2: CAD reconstruction of traffic accident scene (courtesy of [2]).

As *iWitness* was primarily designed for AR and forensic measurement, it 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 modelling packages, especially with CAD systems from CAD Zone [1]. The graphical user interface of *iWitness* is illustrated in Fig. 3, which shows the vehicle collision survey from which the ‘diagramming’ shown in Fig. 1 was produced. *iWitness* has many features over and above the orientation and calibration developments that are to be discussed here. These include fully automatic initiation of all computational functions and automatic recognition of the camera(s) via information contained within the EXIF header of the JPEG or TIFF images. 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 indicates any subsequent improvement or degradation in the spatial intersection accuracy as this review process is undertaken. This provides an effective error detection and correction capability. *iWitness* also supports a centroiding feature which facilitates semi-automatic image point measurement of artificial targets, and even some natural targets, to an accuracy of up to 0.03 pixels.

3. NETWORK GEOMETRY IN AR

As can be imagined, feature points of interest in an AR survey tend to be near planar in their distribution, since the majority lie on or near the road surface. A traffic accident scene can be 50-100m 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 hardly constitute a favourable geometric configuration for close-range photogrammetry. The problem is aggravated by the fact that the camera stations also lie close to the average plane of the object target array. This is well illustrated in Fig. 4, which is both a real and generally representative AR network. When one looks at the plan view, Fig. 4a, the photogrammetric response is that the multi-image geometry is not optimal by any means, but is reasonable. A look at the side elevation plot, Fig 4b, produces a more emphatic response: This is very unfavourable camera station geometry from which to build an initial relative orientation (irrespective of the chosen image pairs) and subsequent multi-image network for bundle adjustment.

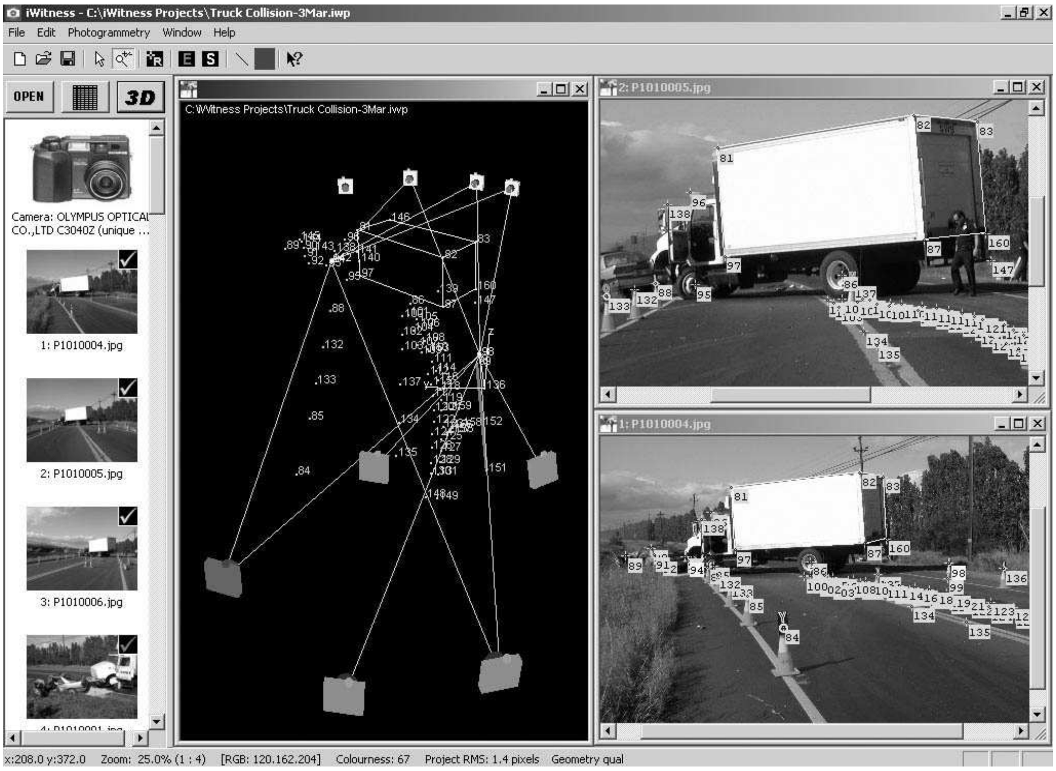

Figure 3: *iWitness* user interface; the CAD diagram in Fig. 1 is from this survey.

However, this is precisely what is required without the aid of any object space control. About the only support to the photogrammetric orientation process is the use of ‘evidence markers’[2], which are back-to-back targets, as illustrated in Fig. 5. These face horizontally and can be semi-automatically measured in *iWitness* via an operated-assisted centroiding function [5]. 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* needed to incorporate some innovative orientation procedures, especially for relative orientation.

4. ROBUST ON-LINE EXTERIOR ORIENTATION

The camera station and object point configuration shown in Fig. 4 illustrates well that photogrammetric network geometry in AR can be complex; far more so in fact 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 may be unsuited to relative orientation and spatial resection.

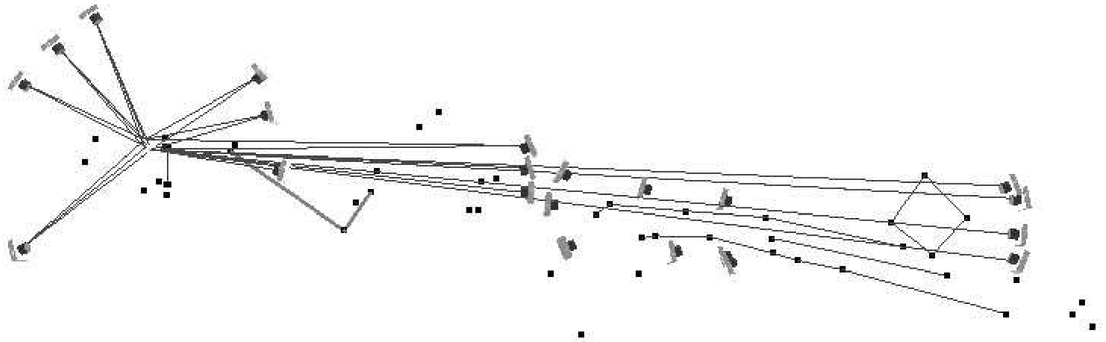

(a) plan view

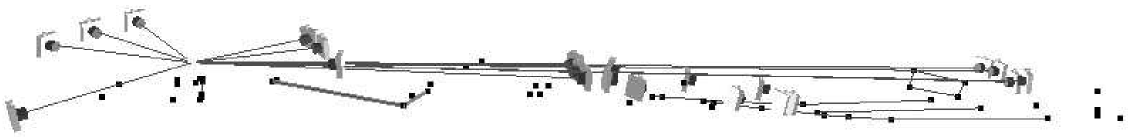

(b) side elevation (tilted for easier interpretation)

Figure 4: Typical near-planar geometry of photogrammetric networks for AR.

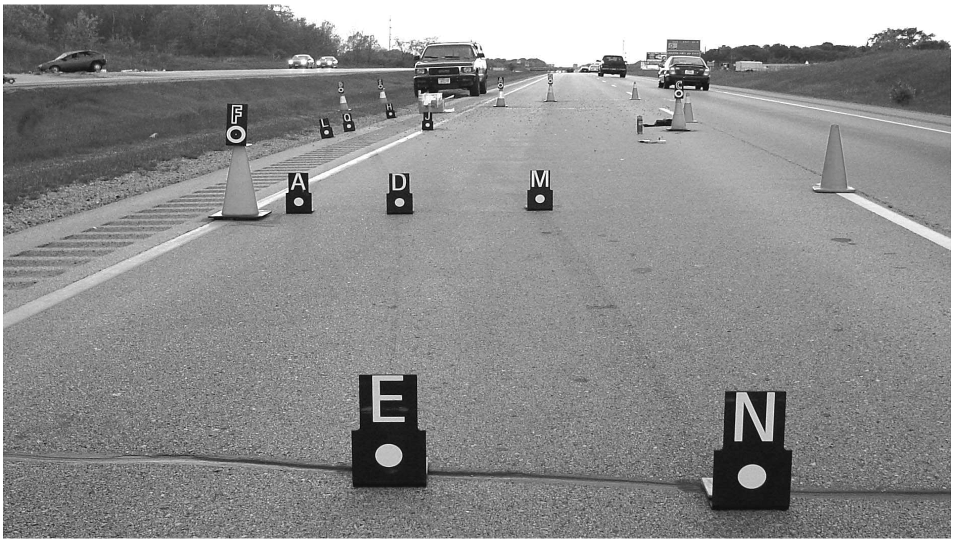

Figure 5: Evidence markers placed on features of interest (photo courtesy of [2]).

Photogrammetrists rely upon two basic mathematical models for sensor orientation: the coplanarity equation for relative orientation, and the collinearity equations for spatial resection, intersection and multi-image bundle adjustment (exterior orientation), with or without camera self-calibration. In their linearized form, both constitute parametric models which are solved via an iterative least-squares adjustment of initial values for the parameters. In the *iWitness* image measurement and orientation paradigm, where the least-squares bundle adjustment is updated as each new observation is made~~,~~ it is imperative that the initial values of the parameters of exterior orientation are determined with sufficient accuracy and reliability to ensure solution convergence.

Traditionally, there have been only two approaches adopted for the determination of preliminary exterior orientation in close-range photogrammetry. 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 then 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 [4,6] as a practical means of providing the necessary 3D control points for automated initial exterior orientation.

A second approach, which has not been widely adopted, is initial relative orientation (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, for example within a stereoscopic model. 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 iterative least-squares 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 the resection problem. The most pressing problem we had in finding 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). Our experience with the least-squares solution to the coplanarity equation is that it is very stable when representative initial parameter values are available, even in situations of very poor geometry.

There has been a wealth of literature within the computer vision community since the Essential Matrix formulation for solving in a linear manner the position and orientation of one camera with respect to another was introduced by Longuet-Higgins [9]. 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 [7]. From reviewing the literature one receives the impression that these approaches had great promise as a means to solve the RO problem. 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. As an aside, we can disregard the fundamental matrix in a photogrammetric context as we always have a reasonable initial interior orientation or 'calibration'. Most consumer-grade digital cameras write the zoom focal length to the EXIF header of the image file and while this does not constitute a photogrammetric principal distance, our experience is that it is generally within 5% of the correct figure.

An evaluation of the essential matrix approach for the estimation of initial RO parameters in *iWitness* was undertaken. Our endeavours, however, were not successful in the context of producing a robust, scene independent RO solution that would be amenable to later refinement via the rigorous coplanarity model. We could immediately discount the prospect of success with near-planar objects, since this is a known failure case – but a geometry that is unfortunately prevalent in AR. We were cautious, however, knowing that either a normalisation process, RANSAC approach or maybe even clever interpretation of the results of a singular value decomposition (and possibly two) could well be necessary to enhance the prospects of success. Also, there were precedents for adoption of the approach in close-range photogrammetry [11], so we persevered – but not for long. In summary, we found the method unreliable and unstable for an application demanding at least a 95% success rate. We also found it unsuited to AR and to the on-line computational scenario utilized in *iWitness*, which seeks to solve the RO as soon as 6 points pairs (8 in the essential matrix case) are referenced. In hindsight we should have taken heed of a comment made by Horn [8]: "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". Or maybe we should have been more suspicious of a method that solves an inherently non-linear problem via a linear model. One can reminisce here on photogrammetric experience with the direct linear transformation.

In our search for a robust procedure for relative orientation in *iWitness* we have settled upon 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, and 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 8 or 9 and occasionally 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 similar to that described in [3]. Generally, the criteria for a correct solution are met after 5 to 6 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.

5. AUTOMATIC CAMERA CALIBRATION

The requirements for camera self-calibration are well recognised: a multi-image, convergent camera station geometry, which incorporates orthogonal camera roll angles, along with an object point array which yields well distributed points throughout the format of the images, and initial starting values for the camera calibration parameters. 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.

With the facility described earlier for robust exterior orientation, forming a self-calibrating bundle adjustment network simply requires the provision of the image point correspondences, ie the (x, y) image coordinates for all matching points. As is now common, the approach to ensuring fast and accurate matching of image point features in *iWitness* is based on coded targets. Novel in the method developed, however, is the use of colour in the codes. Traditionally, codes employed in close-range photogrammetry are geometric arrangements of white dots or shapes on a black background [4]. These geometrically coded targets require optimal exposure to ensure a near binary image is obtained. Such a requirement may be practical for the controlled environments of industrial photogrammetry, but it does not suit the conditions encountered in AR and it does not take advantage of one of the most prominent characteristics of today's digital cameras, namely that they produce colour (RGB) imagery.

The colour codes designed to facilitate fully automatic calibration in *iWitness* are shown in Fig. 6 (albeit without colour due to the greyscale image). Note that the geometric arrangement of the 5-dot pattern is the same; only the colour arrangement varies. Red and green dots are employed to yield 32 (2^5) distinct codes. The blue channel is not utilised in the code approach since the green and red channels yield a far superior response. Once the code dots are detected, a colour transformation process is used to isolate the red/green arrangement and so identify the code. The adoption of colour codes has afforded a more flexible automatic self-calibration procedure.

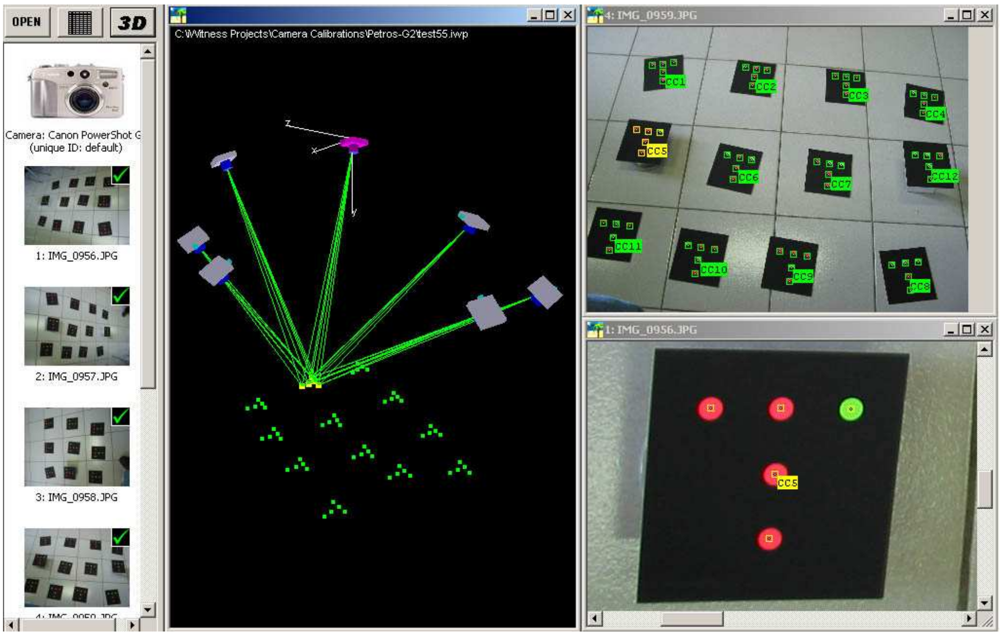

Figure 6. Automatic camera calibration in *iWitness*; note array of 12 colour coded targets.

As for the placement of the codes, it is usually most convenient to simply sit them on the floor, with one or more being out of plane. Non-planarity of codes is not essential for a comprehensive camera calibration, but generally aids in both the initial network orientation, as previously described, and in reducing projective coupling between the interior and exterior orientation parameters. This enhances the precision of the recovered calibration. It has been mentioned that 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. Also shown in Fig. 6 is a typical network for automatic calibration based on colour codes. The codes are purposefully chosen to be relatively large, not to aid in recognition or measurement, but to constitute a sub-group of points. Thus, rather than being treated as a single point, each code forms a bundle of five rays, as is seen in the figure. This means that a broader distribution of image point locations is achieved, which adds strength to the photogrammetric network.

6. CONCLUDING REMARKS

The two innovations described for enhancing the utility, robustness and flexibility of digital close-range photogrammetric systems employing off-the-shelf cameras are incorporated in *iWitness*. Although the development of a new exterior orientation process and an automatic camera calibration strategy utilising 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 1:50,000 of the size of the object, for as little as $2000.

REFERENCES

[1] CAD Zone: http://www.cadzone.com (Web site accessed May 20, 2005).

[2] DeChant Consulting Services – DCS, Inc.: http://www.photomeasure.com (Web site accessed May 20, 2005).

[3] Fischler, M.A. and Bolles, R.C.: Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography, Communications of ACM, 24(6), 381-395, 1981.

[4] Fraser, C.S.: Innovations in Automation for Vision Metrology Systems, Photogrammetric Record, 15(90), 901-911, 1997.

[5] Fraser, C.S and Hanley, H.B.: Developments in Close-Range Photogrammetry for 3D Modelling: the *iWitness* Example, Int. Workshop: Processing and Visualization using High-Resolution Imagery, Pitsanulok, Thailand, 18-20 November, 2004.

[6] Ganci, G. and Hanley, H.: Automation in Videogrammetry, International Archives of Photogrammetry and Remote Sensing, 32(5), 53-58, 1998.

[7] Hartley, R.I. and Zissermann, A.: Multiple View Geometry in Computer Vision, Cambridge Press, 2000.

[8] Horn, B.K.P.: Recovering Baseline and Orientation from Essential Matrix, http://ocw.mit.edu/OcwWeb/Electrical-Engineering-and-Computer-Science/6-801Fall-2004/Readings/, 10 pages, 1990.

[9] Longuet-Higgins, H.C.: A Computer Algorithm for Reconstructing a Scene from Two projections, Nature, 293, 133-135, 1981.

[10] Photometrix: http://www.photometrix.com.au (Web site accessed May 20, 2005).

[11] Roth, G.: Automatic Correspondences for Photogrammetric Model Building, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 35(B5), 713-718, 2004.