

A Feature-Based Matching Approach to Automated Object Reconstruction in Multi-Image Close-Range Photogrammetry

Ida JAZAYERI, Simon CRONK and Clive FRASER, Australia

Key words: Automation, close-range photogrammetry, feature-based matching, interest operators, Poisson Surface Reconstruction

SUMMARY

Automated object surface reconstruction via feature-based matching is commonly employed in topographic and stereo close-range photogrammetry. However, it has rarely been utilised in conjunction with the convergent photogrammetric network geometry that is typically adopted for engineering surveys. This paper describes a new approach to automated 3D object reconstruction from highly convergent, multi-image networks in which the very geometric diversity and redundancy that presents difficulties in area-based matching is a distinct advantage in the feature-based matching approach. Following an initial, automatic exterior orientation of the network, the FAST interest operator is applied to Wallis-filtered images to extract interest points, and a matching of interest points is carried out to yield a dense and accurate 3D point cloud. This is subsequently converted to a triangulated mesh via a Poisson Surface Reconstruction technique, and textured. The paper commences with a brief overview of three important components of the proposed object reconstruction approach: the FAST interest operator, the Wallis filter and the Poisson Surface Reconstruction. Experimental results are then discussed, via an engineering measurement application in which surface modelling is performed to support deformation analysis. The results highlight the practicability, robustness and accuracy of the proposed automated multi-image measurement approach incorporating feature-based matching.

A Feature-Based Matching Approach to Automated Object Reconstruction in Multi-Image Close-Range Photogrammetry

Ida JAZAYERI, Simon CRONK and Clive FRASER, Australia

1. INTRODUCTION

Automated object surface reconstruction from convergent, multi-image close-range photogrammetric networks is commonly employed in high-precision large-scale industrial metrology and engineering surveys. To achieve high-precision, very accurate image point correspondence is required, and this has meant that distinct, artificial targetting is employed, with the targets nowadays generally being either projected onto the object or made of high-contrast, retroreflective material. In addition to the fact that distinct targetted points are of primary interest in many measurement applications, the accuracy and geometric constraints imposed upon image matching within stereo photogrammetry mostly preclude the application of area-based image matching approaches for surface reconstruction of unsignalised objects. Within topographic and stereo close-range photogrammetry, feature-based matching offers an alternative approach to area-based matching via cross correlation or least-squares matching, yet this technique has rarely been employed in conjunction with the highly convergent photogrammetric network geometry that is commonly utilised for engineering surveys.

The creation of 3D models using photogrammetric techniques involves the following phases: image pre-processing, camera calibration and network orientation, image scanning for point detection, surface measurement and triangulation, blunder detection and statistical filtering, mesh generation and texturing, and visualization and analysis. The topic of this paper concerns development of an automated 3D surface measurement strategy suited to multi-image, convergent photogrammetric networks, and that is suited to subsequent high-density mesh generation and texturing. Within the proposed approach to 3D object reconstruction in multi-image networks, the very geometric diversity (highly convergent imagery) and redundancy (multi-ray intersection) that presents difficulties in area-based matching is a distinct advantage in the feature-based matching approach. In the following sections, these aspects are discussed and the results of experimental application of an automated 3D modelling approach incorporating these processes are presented to demonstrate the feasibility of the concept.

Shown in Fig.1 is a representative example of a multi-image, convergent photogrammetric network geometry for an object surface reconstruction survey of the type envisaged here. It is noteworthy that the network geometry renders application of area-based matching impractical because of the excessive perspective disparity and the fact that radiometric similarity cannot be assumed for corresponding feature points. The aim is to recover the surface shape to high-resolution and high accuracy, using multi-ray spatial intersection for extracted and matched unsignalised interest points, even in situation where the object shape is relatively complex.

As a starting point to the proposed feature-based matching approach there has been acceptance of the necessary prerequisite of having a precise exterior orientation for the multi-image network. This is achieved via coded targets, as is common in industrial vision metrology (e.g. Fraser, 2006; Cronk et al., 2006). Thus, in order to achieve the research goal of automated object reconstruction and modelling, attention needs to focus upon two areas: the image scanning/measurement stage, which comprises image pre-processing and interest point detection to provide the features to be matched, and upon the mesh generation for the measured object surfaces. These issues are addressed in the following sections. First, the FAST interest operator, which has been found to be optimal for feature point detection, is overviewed. Second, pre-processing via the Wallis filter is described, since this step has been found to enhance the performance of the FAST operator and subsequent feature-based matching. Thirdly, a Poisson surface reconstruction approach for mesh generation is described. Integration of these concepts into an automated close-range photogrammetric orientation system that utilises coded targets, and possibly uncoded signalised points as well, has enabled realization of a fully automatic object reconstruction and modelling system. Application of the system will be highlighted, with the chosen example involving surface deformation monitoring.




Figure 1: Example network geometry for object surface measurement.

2. THE FAST INTEREST POINT OPERATOR

Interest operators were first developed in the 1970's and since then a number of different algorithms have been developed and presented. For more complete accounts of the range of interest operators, the reader is referred to Schmid et al. (2000), Rodehorst & Koschan (2006) and Remondino (2006). Interest operators detect features of interest in an image, such as corners, edges or regions, and in photogrammetric object reconstruction they are employed to find interest points for matching across multiple images. High quality interest points are required as a preliminary step in this surface measurement process for generating high-accuracy 3D models.

The FAST (Features from Accelerated Segment Test) algorithm, developed by Rosten and Drummond (2006), is a high speed feature detector with strong repeatability properties suited to real-time frame-rate applications, which can perform interest point detection at live PAL

video full-frame rate using less than 7% of the available processing time. A further attribute of the FAST operator, which is of critical importance for image-based modelling from convergent networks, is its invariance to rotation and changes in scale. This represents better performance than many preceding algorithms, including the SIFT operator (Lowe, 2004).

The FAST algorithm is similar in operation to the more familiar SUSAN algorithm developed by Smith and Brady (1997), in that the detector examines a small patch in an image and assesses whether or not it ‘looks’ like a corner. A circular window is scanned across the image and the intensity values of the pixels within or around the window are compared to that of the central pixel. The algorithm considers a circle of 16 pixels around the corner candidate p , as illustrated in Fig. 2, and an interest point is indicated when a set of n contiguous pixels in the circle are all brighter than the candidate pixel I_p plus a threshold t , or all darker than $I_p \leq t$.




Figure 2: 16-pixel circle used by FAST operator.

For each location on the circle $x \in \{1..16\}$, the pixel at that location relative to p (denoted $p \rightarrow x$) can have one of the three states:

$$S_{p \rightarrow x} = \begin{cases} d, & I_{p \rightarrow x} \leq I_p \leq t & \text{(darker)} \\ s, & I_p \leq t < I_{p \rightarrow x} < I_p + t & \text{(similar)} \\ b, & I_p + t < I_{p \rightarrow x} & \text{(brighter)} \end{cases} \quad (1)$$

For each x , $S_{p \rightarrow x}$ is computed for all $p \in P$, the set of all pixels in all training images. This divides P into three subsets P_d , P_s or P_b where each interest point candidate p is assigned a $PS_{p \rightarrow x}$ value. A Boolean variable K_p is then assigned a true value for p being an interest point and a false value otherwise. The algorithm selects an x -value based on the entropy of K_p to provide information about whether the candidate pixel is an interest point. Following determination of the optimal x -value, the process is applied recursively on all three subsets and it only terminates when the entropy of a subset is zero. This means that when all p values in the subset have the same value as K_p , they are all either interest points or non-interest points. A decision tree that classifies all detected points is created from the output of this process, and this is then converted into computer code which is compiled twice for optimisation and used as a corner detector.

The final computation stage of the FAST algorithm involves computation of a corner response function, where non-maximal suppression can be applied, setting a quality measure directly to each candidate corner. The FAST operator does not directly compute a corner response

function and therefore a non-maximal suppression cannot be applied to the resulting features. Instead, a score function V is computed for each detected interest point and non-maximal suppression is applied to remove points that have an adjacent point with a higher V -value. The score function is the sum of the absolute difference between the pixels in the contiguous arc, the centre pixel being given by:

$$V = \max \left(\sum_{x \in S_{bright}} |I_{p \rightarrow x} - I_p| \leq t, \sum_{x \in S_{dark}} |I_p - I_{p \rightarrow x}| \leq t \right) \quad (2)$$

The V score can be used as a quality control measure in post-processing to remove interest points below a chosen threshold. Higher quality interest points, ie points with the highest V -value or absolute difference between pixel intensities in the contiguous arc and the centre pixel, are retained. Jazayeri et al. (2009) have reported upon the performance of the FAST operator, in comparison to the SUSAN and often adopted Foerstner operators.

3. WALLIS FILTER

Enhancement of the images forming the convergent photogrammetric network through pre-processing is often warranted for subsequent feature extraction and image matching (e.g. Baltasvias, 1991; Baltasvias et al., 1996). In the authors' experience, this is indeed the case when utilising the FAST operator. In the proposed approach, the Wallis filter (Wallis, 1976) is applied for this purpose, since studies have shown that interest operators typically find more suitable points on imagery that has been pre-processed with this filter (eg Remondino, 2006; Ohdake & Chikatsu, 2005; Seiz et al., 2002). The Wallis algorithm is adaptive and adjusts pixel brightness values in local areas only, as opposed to a global contrast filter, which applies the same level of contrast throughout an entire image. The resulting image contains greater detail in both low and high level contrast regions concurrently, ensuring that good local enhancement is achieved throughout the entire image. As a result of testing a number of smoothing filters, it was found that the Wallis filter is a most suitable choice considering its ability to provide greater detail in shadowed areas and saturated areas simultaneously, thus allowing a greater number of interest points to be detected. Further details on the Wallis filter in the context of the present investigation are provided in Jazayeri et al. (2009).

4. POISSON SURFACE RECONSTRUCTION

Following the final photogrammetric triangulation stage, a dense, unstructured cloud of 3D points is determined on the object surface(s). For the purposes of 3D modelling and later texturing and visualization, a 3D mesh needs to be generated from the point cloud. A Poisson Surface Reconstruction technique, developed by Kazhdan et al. (2006), has been adopted for mesh generation. The technique is a novel approach that expresses surface reconstruction as the solution to a Poisson equation. Kazhdan et al. (2006) demonstrate that the Poisson algorithm can facilitate the reconstruction of surfaces with greater detail than previously achievable.

The adopted algorithm, which employs an implicit function framework, computes a 3D indicator function; that is, a function that is defined as ‘1’ at points inside the model and ‘0’ outside. It then obtains the reconstructed surface by extracting the isosurface. The algorithm is based on an observation that there is an integral relationship between oriented points sampled from the surface of a model and the indicator function of the model. More specifically, the gradient of the indicator function is a vector field that is zero almost everywhere, except at points near the surface, where it is equal to the inward surface normal. A relationship between the gradient of the indicator function and an integral of the surface normal field is derived in order to compute the vector field \vec{V} of the oriented points. An explicit computation of the vector field would result in a vector field with unbounded values at the surface boundary, so here the indicator function is implicitly solved by convolving the function by a smoothing filter and then considering the gradient of the smoothed function.

Kazhdan et al. (2006) have transformed the computation of the indicator function into a standard Poisson problem. The scalar function χ , namely the indicator function, whose Laplacian (divergence of gradient) equals the divergence of the vector field \vec{V} , is computed:

$$\Delta\chi \equiv \nabla \cdot \nabla\chi = \nabla \cdot \vec{V} \quad (3)$$

The input data \mathcal{S} is a set of samples $s \in \mathcal{S}$, where each sample consists of a point $s \cdot p$ and an inward-facing normal $s \cdot \vec{N}$. Each sample is assumed to lie on or near the surface ∂M of an unknown model M . The input set of oriented points provides precisely enough information to approximate the surface integral with a discrete summation. The input point set \mathcal{P} is used to partition ∂M into distinct patches $P_s \subset \partial M$, and to approximate the surface integral over a patch P_s by the value at point sample $s \cdot p$, scaled by the area of the patch:

$$\begin{aligned} \nabla(\chi_M * F)(q) &= \sum_{s \in \mathcal{S}} \int_{P_s} F_p(q) \vec{N}_{\partial M}(p) dp \\ &\approx \sum_{s \in \mathcal{S}} |P_s| F_{s \cdot p}(q) s \cdot \vec{N} \equiv \vec{V}(q) \end{aligned} \quad (4)$$

Here $\tilde{F}(q)$ is a smoothing filter and $\vec{N}_{\partial M}(p)$ is the inward surface normal at $p \in \partial M$. Following determination of the vector field \vec{V} , the Poisson Surface Reconstruction can solve for the indicator function $\tilde{\chi}$ such that $\nabla\tilde{\chi} = \vec{V}$. However since \vec{V} is not integrable, the algorithm cannot find a direct and explicit solution, and instead adopts a least squares solution, after which the divergence operator is applied to form the standard Poisson equation:

$$\Delta\tilde{\chi} = \nabla \cdot \vec{V} \quad (5)$$

An advantage of the Poisson Surface Reconstruction is that it can be extended to reconstruct non-uniform samples. This is of particular importance to this research as the point clouds generated from the FAST operator are sporadic. The approach in Kazhdan et al. (2006) is to estimate the local sampling density and then scale the contribution of each point accordingly. The algorithm uses a weight function, namely a Kernel density estimator to assign a weight to each point. However rather than simply scaling the magnitude of a fixed-width kernel associated with each point, the algorithm additionally adapts the kernel width. This enables the reconstruction to maintain sharp features in areas of dense sampling while also providing a smooth reconstruction of the surface in areas with sparsely sampled points.

There are a number of advantages to formulating surface reconstruction as a Poisson equation. Poisson systems are well known for their resilience in the presence of imperfect data. In addition, the Poisson Surface Reconstruction recovers the global best-fit model that considers all the input data at once, creating very smooth surfaces that robustly approximate noisy data and requires very little or no post-processing. Finally, since the gradient of the implicit function is constrained at all spatial points, the Poisson Surface Reconstruction attends to spurious 'off-surface' points and succeeds in returning a seamless mesh that closely approximates the input data. The Poisson Surface Reconstruction algorithm is implemented in the Computational Geometry Algorithms Library (CGAL, <http://www.cgal.org>).

5. EXPERIMENTAL TESTING PROGRAM

The experimental testing program conducted to evaluate the multi-ray feature-based matching approach to close-range photogrammetric object reconstruction comprised three phases. The first assessed both the Wallis filter and the degree to which it enhanced interest operator performance, the second evaluated the FAST operator for high-accuracy photogrammetric object reconstruction, and the third assessed the Poisson Surface Reconstruction approach for high-accuracy 3D mesh generation. The object used in the experimental testing was a deformed aluminium plate with graph paper affixed to the surface for quantitative assessment, as illustrated in Fig. 3. These plates are used as surface protection layers in various engineering applications. The plate was then imaged using the convergent network arrangement shown in Fig. 1. Coded targets were placed around the object to facilitate fully automatic network orientation and self-calibration. The resulting exterior orientation was then utilised to determine a dense array of 3D surface points through a feature-based matching of extracted interest points which centred on multi-image point correspondence determination (e.g. Otepka et al., 2002; Sabel, 1999).



Figure 3: Damaged aluminium plate to be modelled for deformation monitoring.

The Wallis filter was first applied to the images of the aluminium plate, and the FAST operator was then run on both the Wallis filtered and original images to ascertain if superior results are obtained when an image enhancement algorithm is applied. For the second phase of testing, the FAST operator was run with different filtering parameters set for the algorithm in order to determine which values would result in the highest accuracy solution with a sufficient number of 3D points. Finally, the Poisson Surface Reconstruction was applied to the resulting point cloud to evaluate its applicability for generation of a high-definition wireframe for the aluminium plate surface.

6. RESULTS

6.1 Image Enhancement Results

The results from the first phase of the experiment indicated that the Wallis filter is indeed a necessary pre-processing function that enables the FAST interest operator to find a greater number of suitable interest points. By applying the Wallis filter, the shadowed areas are brightened and local enhancement is achieved throughout the entire image, as illustrated in Fig. 4a and further discussed in Jazayeri et al. (2009). The result of applying the filter is a normalised image, where the interest operator is able to detect suitable corresponding points in all areas. Issues arising from changes in contrast and illumination are overcome, leading to more repeatable and reliable results.

6.2 Interest Operator Results

The impact of applying the Wallis filter was further assessed in the second phase of the experimental testing, where the results showed that the FAST operator detected on average seven times more interest points on pre-processed images, as highlighted in Figs. 4a to 4c. The results also clearly indicated that the FAST operator is both a very fast and robust algorithm and it yields good localization (positional accuracy) and high point detection reliability, as illustrated both in Fig. 5 and by the results presented in Tables 1 and 2. Table 1 shows the speed and detection rate of the FAST operator when applied to both the original and filtered images. The computation time of the algorithm is a fraction of a second for both images, with over 30,000 interest points being found in the original image and over 200,000 points being detected in 0.3 sec in the Wallis filtered image. The algorithm found points with excellent localization and very few erroneous points were detected. However, to better assess the results of the FAST operator and to obtain a more quantitative measure of its performance, one grid patch on the surface was investigated in more detail.

The results are summarized in Table 2, which shows the total number of points detected, points correctly detected, points missed and points incorrectly detected relative to the total number of points in the selected 11x11 grid patch, in which each node should be detected as an

interest point. These values are averages of results found for all the images forming the network. Here the FAST operator finds nearly all the nodes on the graph paper (88.4% true interest points) with very few erroneous points being detected (1.8% of total points), and with a low number of points missed (11.6% of the total). The table also shows the localization accuracy of the FAST operator, namely how well it positions each interest point relative to its true location. Using a camera with a pixel size of 0.006mm, such as the Nikon D200 used in this test, the localization of the FAST operator would be 3 to 6 μ m and thus clearly suitable for high-accuracy measurement. Fig. 5 shows an enlargement of one of the images to graphically illustrate the results of the FAST operator. The localization accuracy is particularly apparent as the central node is detected at its exact location. Furthermore, there are clearly sufficient points detected by the FAST algorithm.




Figure 4: (a) Wallis filtered image; (b) Enlarged area showing FAST interest operator results on Wallis filtered image; (c) Enlarged area showing FAST operator results on original image.




Figure 5: Enlarged area of the aluminium plate test showing FAST operator results.

Table 1: Speed and detection rate of the FAST operator applied to the aluminium plate.

Original Image		Wallis Filtered Image	
Speed (s)	Detection Rate (#pts)	Speed (s)	Detection Rate (#pts)
0.1	30578	0.3	208963

Table 2: FAST operator results for a selected 11x11, 121 point grid patch.

Total points detected	109 (90.1%)
Correct points detected	107 (88.4%)
Points missed	14 (11.6%)

Incorrect points detected	2 (1.8%)
Localization (pixels)	0.5 – 1
Localization (mm)	0.0031 - 0.0061

In order to have more control over the number of interest points detected by the FAST operator, a filtering function was added to the algorithm. The filter works by assessing the quality of each interest point found, and eliminating all points below a user-defined quality threshold. The quality measure is based on the score function value V (Eq. 2), where interest points with high score values are regarded as being of high quality. As the score function values vary with different images, and since the user would not know which values for V constitute high scores, the filtering function works as a percentage filter. This allows the user to determine what percentages of points are to be retained, based on the score function values of all the points detected, ie if the user selects a filter value of 80%, only 20% of detected points, those with the highest score function values, will be retained. This ensures that only interest points of optimal quality are used in subsequent feature-based matching.

6.3 Object Point Accuracy

Shown in Tables 3a and 3b are summaries of the object point determination results from the multi-image feature-based matching and photogrammetric triangulation of interest points detected via the FAST operator. For each of four Quality Filter values, 80%, 90%, 95% and 99%, the tables list the RMS value of image coordinate residuals, the mean standard error of object point coordinates, the corresponding relative accuracy, the number of points resulting from the final bundle adjustment and the number of erroneous points. The erroneous points comprise those which met image matching criteria but were rejected in the final bundle adjustment. In all cases, the minimum number of imaging rays for an object point was set at 4, the maximum possible being 8.

Table 3a: Object point accuracy for network of original (unfiltered) images.

Quality filter value (%)	RMS of image coord residuals (pixels)	Mean std. error of 3D surface points (mm)	Relative object point accuracy	Number of 3D matched points from final bundle adjustment	Number of erroneous points
80	0.45	0.044	1:18,000	1518	36
90	0.33	0.040	1:20,000	895	4
95	0.28	0.038	1:21,000	504	0
99	0.28	0.029	1:26,000	193	0

Table 3b: Object point accuracy for network of Wallis filtered images.

Quality filter value (%)	RMS of image coord residuals (pixels)	Mean std. error of 3D surface points (mm)	Relative object point accuracy	Number of 3D matched points from final bundle adjustment	Number of erroneous points
80	0.79	0.077	1:10000	6798	3051

90	0.43	0.051	1:16000	3826	177
95	0.29	0.048	1:17000	2278	6
99	0.30	0.038	1:21000	248	0

In applying feature-based matching to the eight images of the plate network, without any quality filtering, some 64,000 points were found to meet the criteria of acceptable matches for subsequent 3D determination. The minimum number of rays set for this matching was three. The results, as anticipated were not acceptable, with possibly 60% or more of the 3D points constituting gross errors. When the FAST quality filter was applied, the number of resulting surface points dropped quite dramatically, as did the number of erroneous points. It can be seen from Tables 3a and 3b, that retention of only 20% of the candidate 2D FAST-detected interest points results in 1518 3D points in the original, unfiltered images, with only 36 of these being wrong point solutions. The corresponding figures for the Wallis filtered case are 6800 and 3050, ie 50% of the points were still erroneous. However, when the quality factor was set to 0.95, the resulting number of valid surface points for the Wallis filter case fell to a still quite dense coverage of 2280, while the number of rejected points dropped to less than 10. The corresponding number of surface points for the unfiltered case is lower, at 500, but here there are no erroneous 3D point determinations. The retention of only the best 5% or so of the detected interest points produced a very reliable 3D point cloud, as judged by the absence of point rejections in the final bundle adjustment process.

As the quality tolerance on the FAST interest points tightens, so the RMS value of image coordinate misclosures in the final bundle adjustment reduces and the internal photogrammetric accuracy thus improves, again as expected. This has a scaling effect on the *a posteriori* standard error estimates for the 3D surface points, which in the case considered reached a level of between 0.03 and 0.05mm for the best 5% of matched interest points. The corresponding proportional accuracy range for the 0.8m diameter aluminium plate segment were 1:16,000 to 1:26,000. The majority of this variation was accounted for by variations in the overall image misclosure values (RMS of image coordinates) rather than by differences in ray intersection geometry. For the 5% of retained interest points, the image coordinate residuals again suggested a matching accuracy in image space of about 0.3 pixels.

Although higher accuracy can be achieved when more points are filtered, the resulting point cloud can become too sparse to meet the requirements for fine resolution mesh generation. On the other hand, if less than 90% of the points are filtered, too many erroneous points may result, thus reducing the final accuracy of the network. This is highlighted by the 80% quality filter case in Table 3(b) for which almost half of the 3D points determined from the Wallis filtered images are rejected as erroneous matches.

6.4 Poisson Surface Reconstruction Results

The mesh generation results indicated that Poisson Surface Reconstruction is a very fast and effective solution for fine resolution, high-accuracy mesh generation. In the case of the aluminium plate measurement test, the surface mesh (Fig. 6a) was obtained fully automatically, with no post-processing required. The results from other test objects also

indicated that the Poisson Surface Reconstruction is well suited to the task of fully automatic mesh generation for accurate 3D surface modelling and representation.




Figure 6: (a) Poisson Surface Reconstruction results, (b) Texturing results for isometric view of plate section and (c) Texturing results for top-down view of plate section.

For many applications, such as surface deformation modelling, generation of the surface model represents the end point of the overall 3D measurement process. In the context of visualisation, however, the triangulated surface model forms the input data to follow-on automatic rendering and photo texturing operations. Discounting issues of radiometric balancing, the texturing is a relatively straightforward operation achieved via back-projection of the mesh triangles into the most appropriate images of the network, these being selected using view-dependence criteria and radiometric considerations. The resulting photo-realistic model of the aluminium plate example is shown in Figs. 6b and 6c.

CONCLUDING REMARKS

In its examination of three operational stages required for the goal of fully automatic 3D object reconstruction from convergent multi-image photogrammetric networks to be realised, this investigation has highlighted the benefits of the Wallis filter for image pre-processing and the applicability of the FAST operator as an ideal interest point detector for feature-based matching. The performance of the FAST operator is very impressive in terms of the number of interest points detected, speed of detection and the accuracy of localisation. The Poisson

Surface Reconstruction results also indicated that the algorithm is well suited to the task of mesh generation for accurate 3D surface modelling and representation. Finally, it is noteworthy that the final surface representations in Fig. 6 were generated fully automatically from image mensuration, through network orientation, feature-based matching and object point triangulation, to mesh generation and texturing, thus illustrating that the viability of the concept of a feature-based matching approach to high-accuracy, detailed surface reconstruction and modelling within convergent close-range photogrammetric networks.

REFERENCES

- Baltsavias, E. P., 1991. Multiphoto geometrically constrained matching, *PhD dissertation*, Institute of Geodesy and Photogrammetry, ETH Zurich, No. 49, 221 p.
- Baltsavias, E.P., Li, H., Mason, S., Stefanidis, A. & Sinning, M., 1996. Comparison of two digital photogrammetric systems with emphasis on DTM generation: Case study glacier measurement, *Int. Arch. Photogramm., Rem. Sens.*, Vienna, 31(B4): 104-109.
- Cronk, S., Fraser, C.S. and Hanley, H.B., 2006. Automatic Calibration of Colour Digital Cameras. *Photogrammetric Record*, 21(116): 355-372.
- Delaunay, B., 1934. Sur la sphere vide. *Bull. Acad. Science USSR VII: Class. Sci. Mat. Nat.*, pp. 793-800.
- Fraser, C.S., 2006. Evolution of Network Orientation Procedures. *Int. Arch. Photogramm., Rem. Sens. & Spatial Inform.*, Dresden, 35(5): 114-120.
- Jazayeri, I., Fraser, C.S. (2009) Interest Operators for Feature-Based Matching in Close-Range Photogrammetry. *Photogrammetric Record* (in press).
- Kazhdan, M., Bolitho, M., Hoppe, H., 2006. Poisson Surface Reconstruction. *Eurographics SGP*, pp 61-70
- Ohdake, T. & Chikatsu, H., 2005. 3D modelling of high relief sculpture using image based integrated measurement system. *Int. Arch. Photogramm., Rem. Sens. & Spatial Inform. Sc.*, 36(5/W17), unpaginated CD-ROM.
- Otepka, J.O., Hanley, H.B. and Fraser, C.S. , 2002. Algorithm Developments for Automated Off-Line Vision Metrology. *Int. Arch. Photogramm., Rem. Sens.*, 34(5): 60-67.
- Lowe, D.G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2): 91-110.
- Remondino, F., 2006a. Image-based modelling for object and human reconstruction, *Ph.D. dissertation*, Inst. of Geodesy and Photogrammetry, ETH Zurich, No. 91, 174p.
- Remondino, F., 2006b. Detectors and descriptors for photogrammetric applications, *Int. Arch. Photogramm., Rem. Sens. & Spatial Inform. Sc.*, 36(3): 49-54.
- Rodehorst, V. & Koshan, A., 2006. Comparison and evaluation of feature point detections. *Proc. of the 5th Int. Symp. Turkish-German Joint Geodetic Days*, Gruendig, L. and Altan, M.O. (Eds.), Technical University of Berlin, ISBN 3-9809030-4-4, Berlin.

- Rosten, E. & Drummond, T., 2006. Machine learning for high-speed corner detection, *Proc. of European Conference on Computer Vision*, Graz, Austria, pp. 430-443.
- Sabel, J.C, 1999. *Calibration and 3D Reconstruction for Multi-Camera Marker Based Motion Measurement*, PhD thesis, Faculty of Applied Physics, Tech. Univ. Delft, Netherlands.
- Schmid, C, Mohr, R & Bauckhage, C., 2000. Evaluation of interest point detectors. *International Journal of Computer Vision*, 37(2):151-172.
- Seiz, G., Baltsavias, E.P.& Grün, A., 2002. Cloud Mapping from Ground: Use of Photogrammetric Methods, *Photogramm. Eng. & Remote Sensing*, 68(9): 941-951.
- Smith, S.M. & Brady, J.M., 1997. SUSAN – A new approach to low level image processing. *International Journal of Computer Vision*, 23(1): 45-78.
- Wallis, K.F., 1976. Seasonal adjustment and relations between variables *Journal of the American Statistical Association*, 69(345):18-31.

CONTACTS

Ms Ida Jazayeri
Department of Geomatics
The University of Melbourne
Melbourne
3010
AUSTRALIA
Tel. + 61 410 760 577
Fax + 61 3 9347 2916
Email: i.jazayeri@pgrad.unimelb.edu.au

Dr Simon Cronk
Department of Geomatics
The University of Melbourne
Melbourne
3010
AUSTRALIA
Tel. + 61 3 8344 0357
Fax + 61 3 9347 2916
Email: cronks@unimelb.edu.au

Prof Clive Fraser
Department of Geomatics
The University of Melbourne
Melbourne
3010
AUSTRALIA
Tel. + 61 3 8344 4117

Fax + 61 3 9347 2916
Email: c.fraser@unimelb.edu.au