

AUTOMATED 3D OBJECT RECONSTRUCTION VIA MULTI-IMAGE CLOSE-RANGE PHOTOGRAMMETRY

I. Jazayeri, C.S. Fraser & S. Cronk

Department of Geomatics, University of Melbourne, VIC 3010 Australia
i.jazayeri@pgrad.unimelb.edu.au, c.fraser@unimelb.edu.au & cronks@unimelb.edu.au

Commission V, WG V/4

KEY WORDS: Object Reconstruction, 3D Modelling, Interest Operators, Feature-Based Matching, Close-Range Photogrammetry

ABSTRACT:

Three important stages within automated 3D object reconstruction via multi-image convergent photogrammetry are image pre-processing, interest point detection for feature-based matching and triangular mesh generation. This paper investigates approaches to each of these. The Wallis filter is initially examined as a candidate image pre-processor to enhance the performance of the FAST interest point operator. The FAST algorithm is then evaluated as a potential means to enhance the speed, robustness and accuracy of interest point detection for subsequent feature-based matching. Finally, the Poisson Surface Reconstruction algorithm for wireframe mesh generation of objects with potentially complex 3D surface geometry is evaluated. The outcomes of the investigation indicate that the Wallis filter, FAST interest operator and Poisson Surface Reconstruction algorithms present distinct benefits in the context of automated image-based object reconstruction. The reported investigation has advanced the development of an automatic procedure for high-accuracy point cloud generation in multi-image networks, where robust orientation and 3D point determination has enabled surface measurement and visualization to be implemented within a single software system.

1. INTRODUCTION

There are two basic approaches to fully automatic object surface reconstruction in close-range photogrammetry. The first employs the artificial targetting of points on the surface, either via markers or projected patterns or points. The second forgoes the use of any surface targetting, relying instead upon image matching to provide a dense 3D point cloud representing the surface. The former approach is at the heart of current commercial vision metrology systems, which are widely employed in large scale metrology. The latter approach, which is overviewed in Remondino et al. (2008), is attractive because it does not require signalisation of the object, but there are a number of constraints associated with it that have thus far precluded its full development. While some limitations, such as automated orientation of networks, are being overcome through the development of new algorithms, others persist. Two of these are the restriction to near 'normal' stereo imaging geometry to facilitate intensity-based matching, and limitations in measurement accuracy and reliability that come as a consequence of these geometric constraints.

What would be ideal in the context of optimising measurement flexibility and accuracy would be the adoption of the automated vision metrology approach of using networks of multiple convergent images, but without the need for targetting. This paper reports on an approach that aims to realise this goal, at least to the extent of avoiding the need to target the object of interest, though the presence of coded targets to support automated orientation remains a requirement. Within the proposed method, feature-based matching is adopted so as to remove constraints upon network geometry. Indeed, the presence of the very convergent imagery that limits intensity-based matching, enhanced the proposed approach.

Image-based modelling has been defined by Remondino (2006) as a complete process that starts with image acquisition and

ends with an interactive 3D virtual model. 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. At present, there are no commercially available systems that allow all of these steps to be executed automatically within the same data processing environment. Instead, the phases are typically completed using multiple software systems that require the 3D data to be translated and interchanged between each of the various measurement, modelling and visualization packages.

The topic of this paper concerns development of an automated 3D surface measurement and visualization strategy, suited to multi-image, convergent photogrammetric networks, that can be incorporated with robust automatic network orientation within a single data processing system. This development centres upon a feature-based matching approach that incorporates image pre-processing via the Wallis Filter (Wallis, 1976), application of the FAST interest operator (Rosten & Drummond, 2006), and adoption of Poisson Surface Reconstruction for 3D mesh generation (Kazhdan et al., 2006).

2. THE FAST INTEREST POINT OPERATOR

Interest operators were first developed in the 1970s 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 power. 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.

The FAST algorithm is similar in operation to the more familiar SUSAN algorithm 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. 1, 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 minus a threshold t .

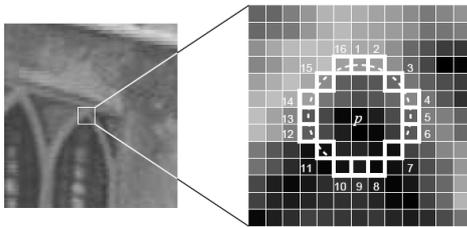


Figure 1: 16-pixel circle used by FAST operator.

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:

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

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 optimization 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 & Fraser (2010) have reported upon the performance of the FAST operator, in comparison to the SUSAN and often adopted Förstner 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 & Fraser (2010).

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.

The problem of computing the indicator function is reduced to finding a scalar function χ , whose gradient best approximates a vector field \vec{V} . Kazhdan et al. (2006) apply a divergence operator to transform this computation of the indicator function into a standard Poisson problem, which is to compute the scalar function χ whose Laplacian (divergence of gradient) equals the divergence of the vector field \vec{V} :

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

The input data S is a set of samples $s \in S$, where each sample consists of a point $s.p$ and an inward-facing normal $s.N$. 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 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 $s.p$, scaled by the area of the patch:

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

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 in the reported approach as the point clouds generated from the FAST operator are sporadic. The approach of 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 require 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. Two test objects, shown in Fig. 2, were used in the testing program. The first was a cultural heritage artefact, namely a replica of a sphinx statue from Persepolis, and the second was a footprint cast subject to forensic analysis.

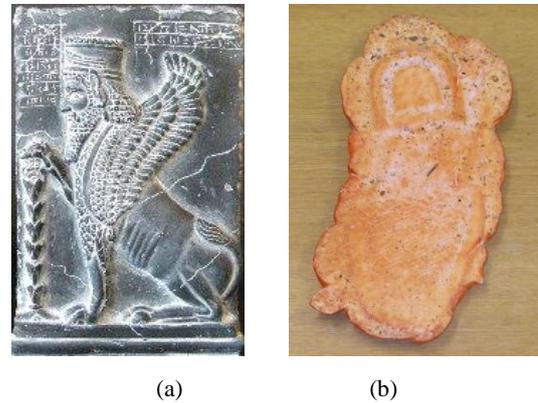


Figure 2: (a) The Sphinx; (b) Footprint cast.

Coded targets were placed around both test objects to facilitate fully automatic network orientation and self-calibration (Cronk et al., 2006), and the subsequent feature-based matching of extracted interest points through multi-image point correspondence determination produced a dense array of 3D surface points. The objects were imaged using a Nikon D200 camera in convergent 8-station photogrammetric networks, the imaging geometry being shown in Fig. 3.

The Wallis filter was first applied to the images of both objects, 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, a quantitative assessment of the FAST operator was performed based on the following performance criteria:

- Speed: A fast and efficient algorithm is required for high-accuracy object reconstruction. It is important that the interest operator has minimal computational time without compromising the accuracy of the results.

- Feature detection rate: This is a measure of the number of true interest points found relative to the number of points missed and points wrongly detected.
- Localization: This measures how well the operator positions each interest point found relative to its true position, which can be assessed visually at pixel level in the image.

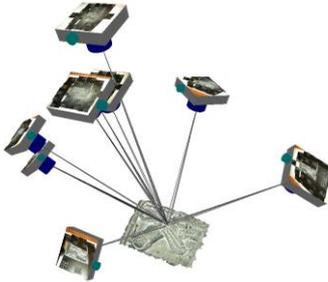


Figure 3: Multi-image convergent network for the Sphinx.

The FAST algorithm was also run with different filtering parameters 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 of each objects 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. 4 and further discussed in Jazayeri & Fraser (2010). The result of applying the filter is a normalized 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. 5a to 5f. 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 by the results presented in Table 1. The table shows the speed and detection rate of the FAST operator when applied to both the original and filtered images of the Sphinx and the footprint cast. The computation time of the algorithm is a fraction of a second for both the original and Wallis filtered images of each test object. More than 10,000 and 13,500 interest points, respectively, were found in the original images of the Sphinx and footprint cast. The corresponding numbers of detected points for the Wallis filtered images were over 75,000 and 150,000. In both cases, the algorithm found points with excellent localization and very few erroneous points were detected.

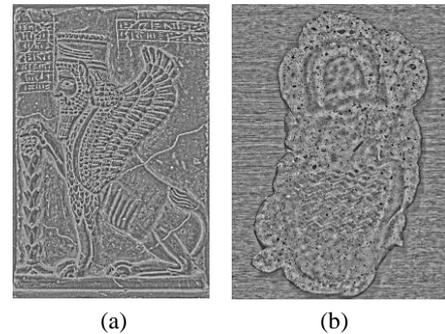


Figure 4: (a) Wallis filtered image of the Sphinx artefact and (b) Wallis filtered image of the footprint cast.

Under the assumption of a localization accuracy of a half to one pixel for the FAST operator, use of a camera with a pixel size of 0.006mm, such as the Nikon D200 used in this test, would yield 3 to 6µm image coordinate accuracy, which is sufficient for automated 3D object measurement. Fig. 5e shows an enlarged section of one of the Sphinx images that contained a symbolic line script with distinct start and end points. The localization accuracy is particularly apparent as these start and end points are detected at their exact locations. Fig. 5f shows an enlargement of one of the footprint images, along with the very dense array of feature points detected by the FAST operator.

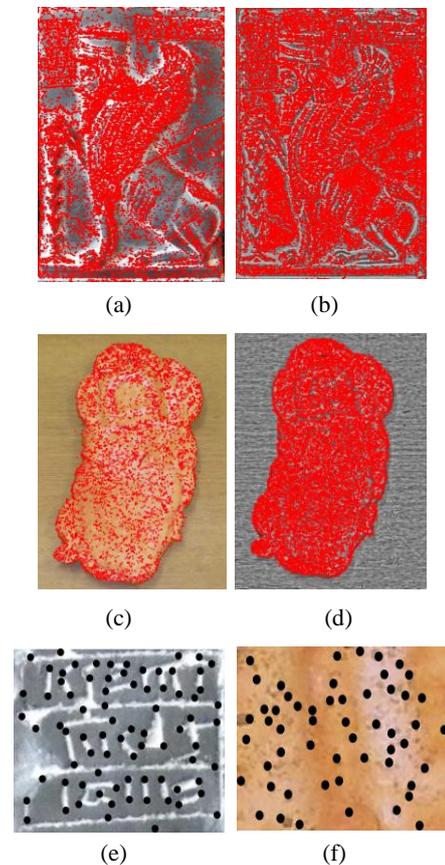


Figure 5: Results of FAST interest operator superimposed on (a) original Sphinx image, (b) Wallis filtered Sphinx image, (c) original footprint image, (d) Wallis filtered footprint image, (e) enlarged area of Sphinx image, and (f) enlarged area of footprint image.

	Original Images		Wallis Images	
	Sphinx	Footprint	Sphinx	Footprint
Speed (sec)	0.1	0.1	0.3	0.4
Detection Rate (pts)	10150	13500	75582	150423
Localization (pixels)	0.5-1	0.5-1	0.5-1	0.5-1

Table 1: FAST operator results.

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, i.e. 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 Feature-Based Matching Results

Shown in Fig. 6 and in Tables 2 and 3 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. Jazayeri & Fraser (2010) report that a Quality Filter value of 90%, ie retention of only the best 10% of interest points, results in high-accuracy 3D points with very few erroneous errors. Tables 2 and 3 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 found at a filtering value of 90%. 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 four, the maximum being eight.

In applying feature-based matching to the eight images of the Sphinx network, without any quality filtering, some 71,000 points were found to meet the criteria of acceptable matches for subsequent 3D determination. The corresponding figure for the footprint network was around 67,000 points. 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 2 and 3, that retention of only 10% of the candidate 2D FAST-detected interest points results in 923 3D points in the original, unfiltered images of the Sphinx, with only 25 of these being wrong point solutions. The footprint cast feature-based matching results yielded 821 3D points with only 22 outliers. The corresponding figures for the Wallis filtered case are 2131 points and 987 outliers for the Sphinx, i.e. 50% of the points were still erroneous. Similarly, the footprint results yielded 1898 3D points with 785 outliers.

The RMS value of image coordinate residuals in the final bundle adjustment, along with the internal photogrammetric accuracy of the Sphinx network suggested a matching accuracy

of about 0.3 pixels for the original images and 0.6 pixels for the Wallis filtered images. The *a posteriori* standard error estimates for the 3D surface points for this network reached a level of between 0.05 and 0.07mm for the best 10% of matched interest points. The proportional accuracy estimates for the Sphinx were 1:14,000 for the original images and 1:10,000 for the Wallis filtered images. The corresponding footprint cast results showed a matching accuracy of about 0.4 pixels in image space, with *a posteriori* object point standard error estimates of 0.047 and 0.07mm for the original and Wallis filtered images, respectively. This translated to proportional accuracy values of 1:18,700 and 1: 15,500.

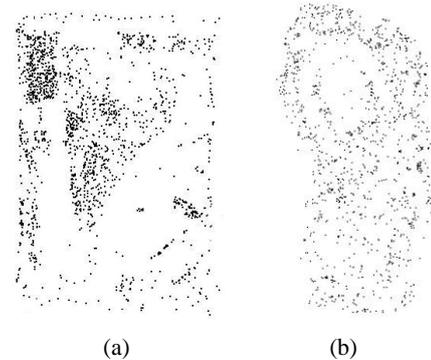


Figure 6: Feature-based matching results for (a) Sphinx and (b) Footprint cast.

	Sphinx	Footprint
RMS of image coord residuals (pixels)	0.29	0.36
Std error of 3D surface pts σ_{xyz} (mm)	0.052	0.047
Relative object point accuracy	1:14000	1:18700
No. of 3D matched pts from bundle adjst.	923	821
No. of erroneous points	25	22

Table 2: Object point accuracy, networks of original images.

	Sphinx	Footprint
RMS of image coord residuals (pixels)	0.61	0.53
Std error of 3D surface pts σ_{xyz} (mm)	0.075	0.070
Relative object point accuracy	1:10000	1:15500
No. of 3D matched pts from bundle adjst.	2131	1898
No. of erroneous points	987	785

Table 3: Object point accuracy, Wallis filtered images.

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 (Jazayeri & Fraser, 2010).

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 Sphinx, the surface mesh, shown in Fig. 7a, was obtained fully automatically, with no post-processing required. The results for the footprint cast 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, as highlighted in Fig. 8a.

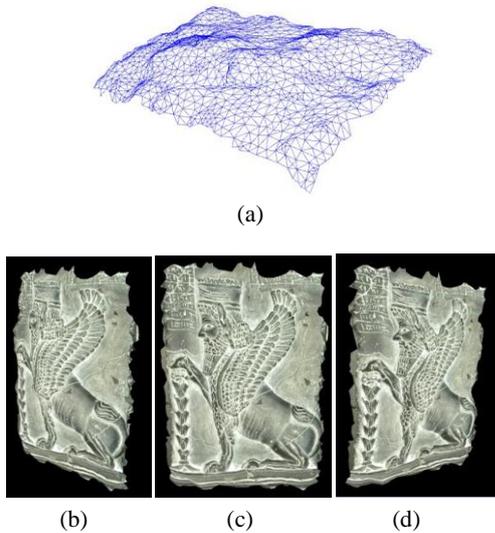


Figure 7: Results for the Sphinx: (a) Poisson Surface Reconstruction, (b) Texturing showing right side of object, (c) Texturing showing front side of object, and (d) Texturing showing left side of object.

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 visualization, 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 both the Sphinx and the footprint cast are shown in Figs. 7b and 7c and Figs. 8b and 8c.

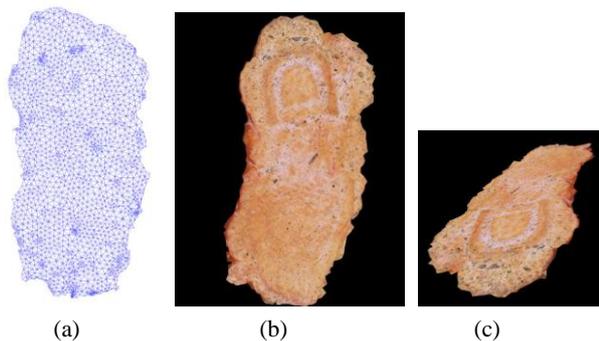


Figure 8: Footprint cast (a) Poisson Surface Reconstruction results; (b) Texturing results showing front of object; (c) Oblique view of object.

7. CONCLUDING REMARKS

In its examination of three operational stages required for fully automatic 3D object reconstruction from convergent multi-image photogrammetric networks, this investigation has highlighted the benefits of the Wallis filter for image pre-processing and the applicability of the FAST operator as an optimal interest point detector for feature-based matching. The

performance of the FAST operator is impressive in terms of the number of interest points detected, speed of detection and accuracy of localization.

The Poisson Surface Reconstruction results also indicated that this 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. 7 and Fig. 8 were generated fully automatically from image mensuration, through network orientation, feature-based matching and object point triangulation, to mesh generation and texturing, thus illustrating the viability of the concept of a feature-based matching approach for high-accuracy, detailed surface reconstruction and modelling within convergent close-range photogrammetric networks.

8. REFERENCES

- Baltsavias, E. P., 1991. Multiphoto geometrically constrained matching, *PhD dissertation*, Institute of Geodesy and Photogrammetry, ETH Zurich, Mitteilungen No. 49, 221 pp.
- 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., Hanley, H., 2006. Automatic calibration of colour digital cameras. *Photogrammetric Record*, 21(116): 355-372.
- Jazayeri, I., Fraser, C.S., 2010. Interest operators for feature-based matching in close-range photogrammetry. *Photogrammetric Record*, 25(129): 24-41.
- 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.
- Remondino, F. 2006. Image-based modelling for object and human reconstruction. *PhD dissertation*, Institute of Geodesy & Photogrammetry, ETH Zurich, Mitteilungen No. 91, 174pp.
- Remondino, F., El-Hakim, S., Gruen, A., Zhang, L., 2008. Development and performance analysis of image matching for detailed surface reconstruction of heritage objects. *IEEE Signal Processing Magazine*, 25(4): 55-65.
- Rodehorst, V., Koshan, A., 2006. Comparison and evaluation of feature point detections. *Proc. 5th Int. Symp. Turkish-German Joint Geodetic Days*, Gruendig, L. and Altan, M.O. (Eds.), Technical University of Berlin, Berlin.
- Rosten, E., Drummond, T., 2006. Machine learning for high-speed corner detection, *European Conference on Computer Vision*, Graz, Austria, pp 430-443.
- 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, *Photogrammetric Engineering & Remote Sensing*, 68(9): 941-951.
- Wallis, K.F. 1976. Seasonal adjustment and relations between variables *Jnl. American Statistical Association*, 69(345): 18-31.