

EVALUATION OF LOW-COST TERRESTRIAL PHOTOGRAMMETRY FOR 3D RECONSTRUCTION OF COMPLEX BUILDINGS

S. Altman^a, W. Xiao^a, B. Grayson^a

^a School of Civil Engineering and Geosciences, Newcastle University, NE1 7RU, Newcastle upon Tyne, UK –
simon88altman@gmail.com, wen.xiao@newcastle.ac.uk

Commission II, WG II/1

KEYWORDS: Low-cost, Photogrammetry, Structure from Motion, Accuracy, Terrestrial, 3D modelling

ABSTRACT:

Terrestrial photogrammetry is an accessible method of 3D digital modelling, and can be done with low-cost consumer grade equipment. Globally there are many undocumented buildings, particularly in the developing world, that could benefit from 3D modelling for documentation, redesign or restoration. Areas with buildings at risk of destruction by natural disaster or war could especially benefit. This study considers a range of variables that affect the quality of photogrammetric results. Different point clouds of the same building are produced with different variables, and they are systematically tested to see how the output was affected. This is done by geometrically comparing them to a laser scanned point cloud of the same building. It finally considers how best results can be achieved for different applications, how to mitigate negative effects, and the limits of this technique.

1. INTRODUCTION

3D modelling is becoming increasingly fundamental to the documentation of the built environment. Laser scanning is the industrial standard for this, enabling fast, reliable and highly accurate capture of point cloud data. However, it remains an expensive technique available only to companies who are able to make the investment. This limits its use, especially in less economically developed parts of the world, making cheaper alternatives a desirable. This is particularly relevant in the field of heritage conservation, where accurate 3D models of buildings are used for maintenance and restoration, recording ornate, irregular features in a way that is hard to match with 2D plans or photographs. For example, the recent earthquake in Nepal damaged or destroyed many historically significant buildings, and digital models would have been greatly beneficial to their restoration (Dhonju *et al.*, 2017).

Compared to laser scanning, terrestrial photogrammetry is cheaper and more portable. At a minimum it requires a consumer-grade digital camera, a computer and some free software. This is widely attainable, even in poorer countries. Results may be strengthened by higher quality DSLR, proprietary software and possibly the addition of control points, which would require a total station. Though more expensive, this still costs less than laser scanning, and would certainly expand the global availability of digital building modelling if adequate. Additionally, if surveying in areas with a high crime risk, a digital camera would be a far smaller loss than a laser scanner; plus its portability means the surveyor can leave the scene much quicker if they sense danger.

UAVs are very useful for photogrammetry. But they can be expensive, difficult to fly and their use is often restricted in urban areas. Terrestrial photogrammetry is potentially useful as it is far less restrictive. Therefore this study will investigate how accurately it can capture complex buildings. There is specific focus on which variables of the process have the greatest effect on the outcome, how much skill does the user need, and how cheaply can acceptable results be

produced. Variables tested will be 1) different software, 2) the photography strategy and 3) the use of control points. Different photogrammetric point clouds of a building will be produced, and geometrically compared to a laser scan of the same building. Results will hopefully demonstrate the accuracy capabilities of this process and provide insight in how to achieve best results.

2. RELATED WORK

Many studies have explored different techniques for low-cost image based 3D documentation of cultural heritage (Boochs *et al.*, 2007; Remondino, 2010). Reu *et al.* (2013) recently tested a low-cost reconstruction approach to archaeological sites using Agisoft Photoscan. Highly-accurate results were produced although it was emphasised that technical knowledge and skill is important. An experiment on modelling ancient tombs in Oman was conducted by Banse *et al.* (2015). They demonstrated that DIM produced point clouds of comparable accuracy to TLS, except in the gaps between the bricks, which were obscured by shadow.

Dhonju *et al.* (2017) evaluated the practicalities and benefits of low-cost image based modelling for heritage preservation in Nepal, following the devastating 2015 earthquake. It concluded that it is of great value considering the many thousands of heritage structures that are in danger worldwide, but tall buildings pose problems which are hard to solve without a UAV or similar.

There are also numerous studies on image based modelling for complex modern buildings. Ground-based photogrammetry was compared to laser scanning for accurate modelling of architectural sites by Gonizzi Barsanti *et al.* (2013). Comparable results were produced although repeatable accuracy from image based modelling was found to be lacking. Ippoliti *et al.* (2015) evaluated the advantages and limitations of image based reconstruction for

architectural surveying. The site was a complex historic courtyard. 3D models and orthographic plans were successfully produced although practitioner expertise was found to be crucial.

Remondino et al. (2012) tested 5 free or low-cost image reconstruction softwares on sites of varying scale. Results showed that software makes a difference and the results are not always consistent. A similar study by Kersten et al. (2015) using various well-known software packages concluded that image based models could not achieve the same accuracy as laser scanning.

Accuracy assessment of various softwares for low-cost UAV image based modelling was carried out by Oniga et al. (2017) on a complex hyperbolic-paraboloid shaped building. Results varied considerably demonstrating the impact of software on airborne image data. The influence of data processing methods on reconstruction from UAV imagery was investigated by Caroti et al. (2015), confirming ground control points are essential for accurate reconstruction of airborne data.

3. DATA ACQUISITION

The test building is The Church of St. Thomas the Martyr in Newcastle Upon Tyne, UK. The 30m tall front section will be captured. It has many ornate features, so detailed photography from the ground is a challenge. This is compounded by the study area being in a busy part of town with many pedestrians, a main road and trees (Figure 1). The road forces data capture from either closer or further than is ideal for such a tall building. Crowds limit where tripods can be placed and find their way into photographs. The trees also obscure view from many places. The scene is far from ideal, but provides a realistic reflection of challenges faced in urban surveying.




Figure 1. Google Earth 3D model of the Church and the surrounding area (Accessed 31/03/2017)

The laser scanning was captured in a local coordinate system from 5 positions using the Leica P40. The laser scan is considered to be 'ground truth', although in reality it

contains some uncertainty which must be considered in the final assessment. The P40 has stated 3D positional accuracy of 3 mm over 50 m. The mean registration error was 3 mm. So the maximum error of the scan is 6 mm.

Photographs were captured with the Canon EOS 600D, with 5184 x 3456 resolution, 18mm focal length and auto-adjust lighting turned off.

Two sets of 150 photos were captured on separate days (Figure 2). The first was taken closer to the building, hence more detail but difficult to capture the higher parts, plus the sun was high in the sky causing some glare issues. Radial distortion was expected to be a problem due to the close proximity and steep upwards angle of some images. For the second set, many of them were taken on the other side of the road. This was a little too far hence detail was lacking, but it was thought they might improve the overall geometry of the model as the entire church was contained near the image centre where radial distortion is lowest. The sun was better positioned but some photos were partly obstructed by trees.




Figure 2. Camera positions of photostet 1(left) and 2 (right)

3.1 Control Measurements

Control points were captured at various points on the building with the Leica TS09 total station. On the far side of the road there were too many viewing obstructions so measurements were taken on the near side. 20 well spread points were measured from 2 stations, although it was too close to tilt the instrument enough to reach the upper turrets. The roundedness of the old stone made the points a little ambiguous, so 8 were measured from both stations for redundancy. These points had a mean 3D positional difference of 3mm.

For scaling the uncontrolled models, a single scale bar measurement was made with a Leica Disto. This could have alternatively been performed with a long tape measure.

4. METHODOLOGY

4.1 3D Reconstruction

This study will firstly involve making a range of 3D reconstructions of the photographs. These will be varied in terms of their input images, software and use of control.

4.1.1 Software

Photogrammetric outputs are heavily dependent on the algorithms used, and these vary between different software. Therefore several will be tested to see how much it affects results.

- *Agisoft Photoscan*: A low-cost, user-friendly software used widely for commercial and research purposes. For commercial reasons, little information is available about the algorithms used. It is the primary software used in this study.

- *Apero / MicMac*: An open-source software duo that uses a modified SIFT++ feature extractor (Vedaldi and Fulkerson, 2010) and a Gauss-Newton bundle adjustment (Pierrot-Deseilligny and Cléry, 2011). It is widely used in research, mainly in command-line format. GUIs have recently been produced which will likely expand its use to less technical users.

- *Bentley ContextCapture*: A user-friendly, proprietary software aimed at professional and academic users. For commercial reasons the algorithms are not publicised.

- *Visual SFM*: A free GUI software aimed at casual users. Their own feature extractor called SiftGPU is used with a multicore bundle adjustment (Wu *et al.*, 2011). Control points cannot be incorporated in the reconstruction.

An initial test will be conducted by aligning uncontrolled models from each software to the laser scanned point cloud. The cloud-to-cloud distances will then be compared visually, and from this, one commercial and one open-source package will then be chosen. They will be used to create controlled models upon which more rigorous statistical tests will then be conducted.

4.1.2 Control Points



Figure 3. Configuration of the 20, 12 and 3 control points (left to right)

Control points are needed to reference the final models. But in theory they should also improve the geometry beyond that achieved by the software's self-calibration procedures. So models from the same photoset will be created with and without control to test whether this is true, and to what extent. The controlled models will have their control point accuracy set to 3mm. It is also tested whether the total number of control points used has a significant effect. This is tested in both Photoscan and MicMac to see if there is any change in results and whether this varies between software. Models will be created from the same photoset with the full 20 points

and also reduced sets of 12 and the minimum which is 3 (Figure 3). The uncontrolled models are scaled in Photoscan using the scale bar function.

4.2 Point Cloud Pre-processing

Cloud Compare is a free, user-friendly software with a multitude of tools for the processing and analysis of point clouds. It was used to assess accuracy of the different photogrammetric point clouds. Every photogrammetric model was cleaned of noise to the same standard so fair statistical tests could be performed. This included removing sky (mostly achieved by masking) and any surrounding objects or pavement. The glass church doors were photogrammetrically reconstructed, but the lasers went through them. Therefore they were removed as they could not be compared properly. Any noise within the church body was left, as it could be considered a reconstruction flaw and hence should be included in the accuracy assessment. Perfect consistency with noise removal was impossible, but variations were small and should barely affect each model's global statistics.

All photogrammetric point clouds are aligned to the laser scan using the Iterative Closet Point (ICP) alignment tool. Whilst some point clouds were created with control points, this was to test the effect on their geometry, not to georeference them, and they are in a different local coordinate system to the laser scan. So they are still aligned by ICP to keep the test consistent. This rotates and translates it to best fit the laser scanned cloud, without changing the scale. Outlying points are excluded from the process in order to improve the fit.

4.3 Accuracy Comparison

The cloud-to-cloud (C2C) distance is the measured between each photogrammetric cloud and the reference laser scanned cloud. This is done by taking every point of the former and finding its nearest neighbour on the latter (Figure 4). This will not be the 'true distance', but since the reference cloud is of sub-millimetre density with little noise, the difference should be negligible. It is possible to convert the reference cloud to a mesh, but the density is such that improvements would be negligible. M3C2 distancing is an alternative tool in CloudCompare which aims to make a smarter assessment of each point's nearest neighbour based on its normal (Lague, 2013). It is better for both positive and negative distance, but beyond the scope of this study in which overall accuracy is the question. The points can be colour coded by their C2C distance value to show which areas match the reference cloud closely and which do not. Points with a C2C distance >80cm were removed from all clouds as they were obviously outliers. The mean scalar value is taken from each cloud as an indication of their overall accuracy. Their standard deviation is also taken as an indicator of noise.

However, this approach does not test the completeness photogrammetric clouds. To account for this, C2C distances are also measured in the other direction, from the laser scan to the photogrammetric clouds. This laser scan cloud is complete, so if there are gaps in the photogrammetric cloud

then the C2C distances around there will be large (Figure 5). More gaps will mean a higher mean C2C distance. These gaps can be shown by colour coding the points, and the overall mean scalar value will give an indication of each cloud's completeness. The standard deviation is not used as it is more relevant when measured in the other direction. Together, these tests give an overall indication of each cloud's quality based on accuracy, noise and completeness.

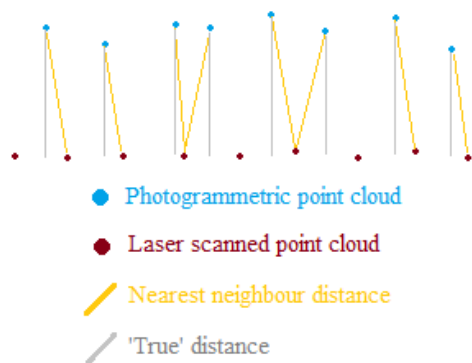


Figure 4. The nearest neighbour distance that cloud-to-cloud distancing is based on




Figure 5. Measuring C2C distances from the laser scanned point cloud will capture gaps in the photogrammetry model

4.4 Automatic Feature Extraction

To explore how useful photogrammetric point clouds are for documentation compared with laser scanning, some extraction tools will be tested on both under the same conditions. 3D ReShaper is a commercial software with a variety of semi-automated tools for point cloud development. Meshes can be automatically created from the points by triangulating them. The smoothness and sharpness of the results give some indication of quality and usability. This is further tested by trying to extract break lines from the mesh. Results will only be visually assessed (not statistically) in order to draw some conclusions about the usefulness of image based reconstruction.

5. RESULTS AND ANALYSIS

5.1 Different Software and Control Points

Uncontrolled models were made from all 4 softwares, with the same 150 photos of the first photoset. C2C distancing showed Photoscan and ContextCapture to give results of

comparable accuracy, low noise but overall Photoscan produced better geometry. MicMac had accurate geometry but appeared noisy with some gaps. VisualSfM gave noisy results with weak geometry (Figure 6). This clearly indicates that different reconstruction algorithms significantly affect the quality of results. Photoscan and Micmac were chosen for further analysis.

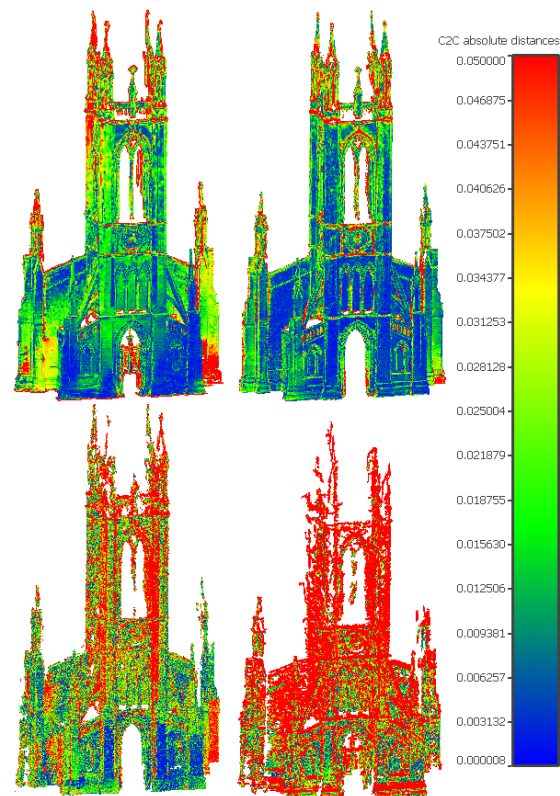


Figure 6. C2C distances of different software. Clockwise from top-left: ContextCapture, Photoscan, VisualSfM, MicMac. Red indicates 5cm and above

In testing different numbers of control points, Photoscan was barely affected. C2C distances were all within 1mm, even with no control points at all (Figure 7). MicMac had lower mean C2C distances overall, with control points clearly making a positive difference. It was also more affected by the number of them, with C2C distance decreasing as number of points decreased. This is surprising as usually more control points will have better accuracy. It is possible that the control points chosen for the reduced sets were coincidentally more accurate than the others thus improving the transformation. Or it is affected by the accuracy of the laser scanned point cloud. Further testing would be required to draw any conclusions. The standard deviations were fairly consistent for both software, indicating that number of control points has little effect on noise levels.

When measuring from the laser scanned clouds, Photoscan had a much lower mean C2C distance, at 22mm to MicMac's 48mm. This is due to more gaps in the MicMac clouds.

Whilst results initially suggest that MicMac is more accurate than Photoscan due to lower mean C2C distances, this does not necessarily mean MicMac is better. The reconstruction is accurate, but there are more gaps, despite those areas being included in the photographs. So it might be that the algorithms demand a greater certainty in the dense image matching, and if an area is below a certain reliability threshold then it is left out. As Photoscan is more commercially focussed, completeness of results might be a greater priority. If so there may be a lower threshold for the areas it will try to reconstruct.

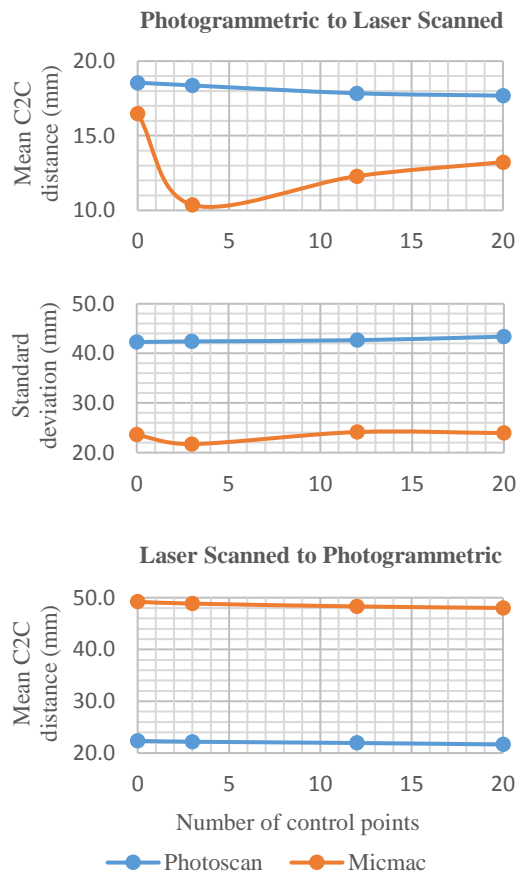


Figure 7. C2C distance testing of point clouds in Photoscan and MicMac

5.2 Number of Photos

Different numbers of photos from the first photoset were used to see if less photos means less accuracy. This was tested with and without control. Additionally, the same control points were again used but this time unconstrained. To achieve this their accuracy description was set much lower at >1000m, to see if this would reduce their effect on the outputs.

Without control, the mean C2C distances were stable as the photoset dropped in number to 90, and then it gradually increased (Figure 8). The difference between 30 and 150 photos was only 4mm. The controlled set had a mean C2C of 1-2mm lower than the unconstrained, which remained

stable until 30, when it increased sharply. The models with unconstrained control had slightly greater mean C2C distances than the standard control, but the difference was never more than 1mm. This indicates that control points do strengthen model geometry, but their level of constraint is of little consequence, at least in Photoscan.

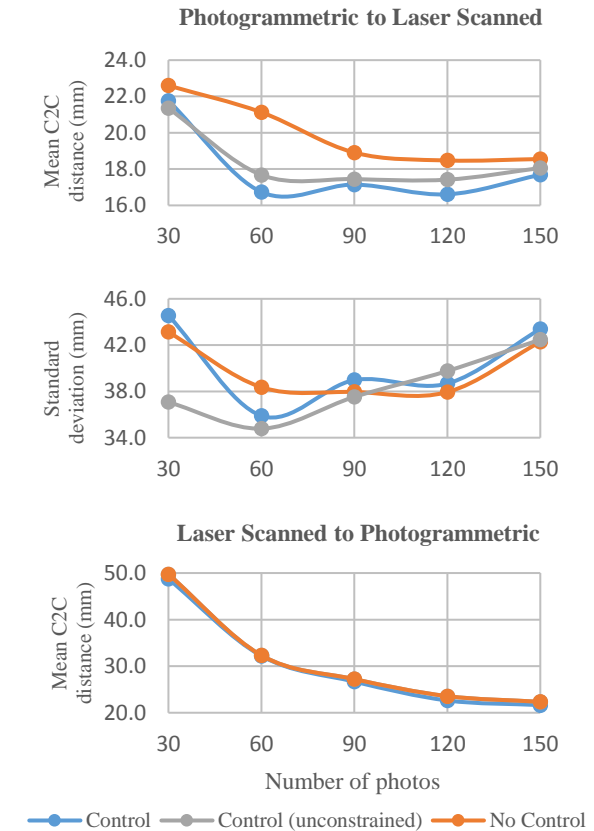


Figure 8. C2C distance testing of photosets with different numbers of images

It can be seen when colour mapping the C2C distances that the inaccuracy is greatest at the upper turrets, particularly on the 30 photo sets (Figure 9). This part of the church is not captured as well so it makes sense that it requires greater redundancy from more photos or accuracy will suffer.

The trend of standard deviation is more complicated. The controlled models from 60 photos are lowest, and from 120 to 150 there is a significant increase across all models. It may suggest that more photos can increase noise and inaccuracy of the model. A possible explanation could be that the photosets with 150 contained more photos with variety of scales and geometries, which interfered with the reconstruction. For the sets with fewer photos, these were the first to be removed. Then both the distance and standard deviation are increasing when there are too few (30) photos.

Measuring from the laser scanned cloud produced results closer to what would be expected. Care was taken to ensure that even the set of 30 photos contained all parts of the church, although overlap may have been weak in places. C2C distance increases as number of photos falls, indicating

that fewer photos means more gaps in the model or the model is less accurate.




Figure 9. C2C distances of Photoscan point clouds with 30 photos (left) and 150 photos (right). Red is 5cm and above

5.3 Photoset Arrangement

The first photoset produced much stronger results than the second. Mean C2C distances were around 1cm lower (17.7mm to 27.0mm) (Figure 10). For photoset 2 it was thought that a mix of photos from close and further away might reduce distortion at the upper turrets, since photos from afar there is a less upward camera angle, and more of the building is towards the lens centre where radial distortion is lowest. The general geometry of the models was good, with few gaps, but they suffered from noise and lack of definition (Figure 11). This is most likely due to some photos being too far away (due to the road), hence lacking detail which caused tie-point matching errors. Extreme scale differences could also be contributing. There were also obstructions in some of the photos (trees, vehicles, pedestrians, etc.), which could also have interfered with this.


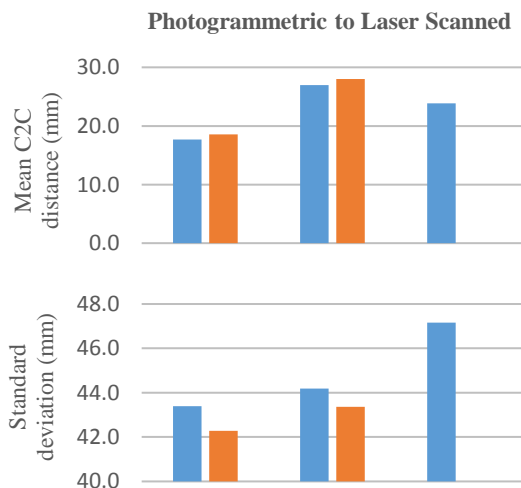


Figure 10. C2C distance testing of models generated from differently arranged photosets

The photosets were combined to see if this would improve the reconstruction due to redundancy. This was not the case. Rather, the accuracy was approximately halfway between the first and second photoset. This would suggest that greater numbers of photos do not necessarily improve the model; if the set contains photos that are not up to standard, then they are more likely to reduce redundancy by causing tie-point ambiguity.




Figure 11. C2C distances of Photoscan point clouds with photoset 1 (left) and 2 (right). Red indicates 5cm and above

5.4 Feature Extraction

A sharp, smooth and complete mesh was extracted from the laser scanned points. The mesh extracted from the Photoscan results was also comparably smooth and complete, but the edges were more rounded. The MicMac mesh had sharper edges, but had gaps and was rougher due to noise.

Visually, the breaklines extracted from the laser scan and from Photoscan were almost indistinguishable. Though far from perfect they could both be useable as a basis for creating 2D vector plans. This could be investigated further in future studies. The breaklines from MicMac were noisier and less complete, and would be harder to use (Figure 12).




Figure 12. Mesh and breakline extraction from point clouds. Top to bottom: Laser scan, Photoscan, MicMac.

6. DISCUSSION

The mean C2C distances of the various reconstructions were all within 1-3cm (Photogrammetric to Laser Scanned). The more successful models had over half their points within 1cm, and over 95% of points within 5cm (Figure 13). The

lowest mean C2C distance was 10.4 mm, achieved by the MicMac cloud with 3 control points. The model with lowest mean C2C distance without control was 16.5 mm, also by MicMac. Photoscan's best with control was 16.6 mm (120 photos) and best without control was 18.5 mm (also 120 photos). Note that the uncertainty of the reference laser scan has not been accounted for.




Figure 13. Histogram of C2C distances for the Photoscan controlled model with 150 photos. 54.7% of points are <1cm and 95.9% are <5cm

Accurate tie-point matching is fundamental to the quality of results, and the photography strategy does seem to have a big effect on this. Increasing the number of photos does improve accuracy of the reconstruction by adding redundancy, but only if they are of the necessary standard. Photos lacking definition or containing obstructions will only serve to introduce ambiguity. The quality of the camera lens will therefore be a factor, and the user must consider how the obstacles surrounding the building will impact the photo positions. Large scale differences within the photoset are also likely to increase ambiguity. Every scene poses different challenges, so experience through trial and error is essential in learning to achieve consistent results.

Taking photos at a steep upward angle is difficult to avoid in terrestrial photogrammetry. It can cause issues with glare, plus the buildings features are hard to identify from such extreme perspectives. Photos can be taken from further away, but image detail is important and must not be sacrificed too much. Therefore with tall buildings, it can be hard not to lose accuracy at the higher parts. This is where UAVs are useful.

The need for control points depends on the intended application of the model. They are essential if geo-referencing is required, but if not then it is a question of how geometrically accurate the output must be. Best fitting a model to a network of control points (by least-squares or some other method) is a more robust method of scaling than using scale bars. The control points also improve accuracy by geometrically correcting the model, but the extent of this depends on the software used.

MicMac seemed to avoid reconstructing areas if they did not meet a high threshold of certainty. This meant the reconstructed areas were accurate, especially if control points were used, but there were more gaps. Photoscan results were highly accurate in parts, but appear to place greater emphasis on reconstructing all areas, even if there is some ambiguity in parts. There were still gaps, especially

when fewer photos were used, but the threshold of certainty at which it will attempt reconstruction seems to be lower.

It is useful to have complete models, and as a commercial software it is no surprise Photoscan tries to deliver a 'finished product'. Results of the feature extraction demonstrate that it does this well. The trade-off however, is that the true accuracy of the model is less certain. The feature extraction results from MicMac are not bad, but show that more work would be needed from the user to complete it. It is no doubt possible to produce complete models in MicMac as well, but it likely requires a more refined image set. Therefore the scene would need to allow for this, plus greater understanding would be needed by the practitioner.

7. CONCLUSIONS

Terrestrial photogrammetry is potentially a very useful tool for modelling complex buildings. But it has its limits and the user must be aware of them. If visually strong results are needed and some centimetre level inaccuracies can be tolerated, it is very useful. For creating digital records and 2D vector plans; some knowledge, a good camera and a commercial software package would mostly give adequate results. If higher accuracy is needed; dedicated practice and experimentation to find the right software, equipment and methods could probably produce sub-centimetre results. Control points are definitely recommended if this is the aim. To achieve reliability, experience will be essential.

There are, however, many variables in terrestrial photogrammetry which are hard to predict or control. Therefore its use for more precise structural analysis might be limited. It would not be able to reliably detect millimetre level building subsidence, for example. But it could have uses in detecting more significant change like earthquake damage or fallen bricks. In these cases, measuring C2C distance from an older model to a current one would reveal what has fallen away. Models of tall buildings will diminish in accuracy towards the top, and this is hard to overcome. Using scaffolding or long poles to raise the camera are some potential low-cost solutions worth exploring. The effects of different camera lenses would also be worth investigating.

REFERENCES

- Banse, P., Kluge, P. and Runne, H. (2015) Comparing DIM to TLS, *GIM International*. iss. 7, vol.29, pp. 27-29.
- Boochs, F., Heinz, G., Huxhagen, U. and Müller, H. (2007) Lowcost image based system for nontechnical experts in cultural heritage documentation and analysis. In *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. pp. 165–170.
- Caroti, G., Martínez-Espejo Zaragoza, I. and Piemonte, A. (2015) Accuracy Assessment in Structure from Motion 3D Reconstruction from UAV-Born Images: The Influence of the Data Processing Methods, *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, vol. XL-1/W4, pp. 103-109.
- Dhonju, H. K., Xiao, W., Sarhosis, V., Mills, J. P., Wilkinson, S., Wang, Z., Thapa, L. and Panday, U. S. (2017) Feasibility Study of Low-cost Image-Based Heritage Documentation in Nepal, *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, vol. XLII-2/W3, pp. 237-242.
- Gonizzi Barsanti, S., Remondino, F. and Visintini, D. (2013) 3D Surveying and Modelling of Archaeological Sites. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, II-5/W1, pp. 145-150.
- Ippoliti, E., Meschini, A. and Sicuranza, F. (2015) Structure from Motion Systems for Architectural Heritage. A Survey of the International Loggia Courtyard of Palazzo Dei Capitani, Ascoli Piceno, Italy. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, XL-5/W4, pp. 53-60.
- Kersten, T., Mechelke, K. and Maziull, L. (2015) 3D model of al zubarah fortress in Qatar—Terrestrial laser scanning vs. dense image matching. *International Archives of the Photogrammetry, Remote Sensing Spatial Information Science*, Volume XL-5/W4, pp. 1-8.
- Lague, D., Brodu, N. and Leroux, J. (2013) Accurate 3D comparison of complex topography with terrestrial laser scanner: Application to the Rangitikei canyon (N-Z), *ISPRS Journal of Photogrammetry and Remote Sensing*, vol.82, pp. 10-26.
- Oniga, E., Chirilă, C. and Stătescu, F. (2017) Accuracy Assessment of a Complex Building 3D Model Reconstructed from Images Acquired with a Low-cost UAS, *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, XLII-2/W3, pp. 551-558.
- Pierrot-Deseilligny, M., Cléry, I. (2011) APERO, an open source bundle adjustment software for automatic calibration and orientation of a set of images. *Int. Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38(5/W16).
- Remondino, F. and Rizzi, A. (2010) Reality-based 3D documentation of natural and cultural heritage sites—techniques, problems, and examples. *Applied Geomatics*, 2(3), pp.85–100.
- Remondino, F., Del Pizzo, S., Kersten, T.P. and Troisi, S. (2012) 'Low-Cost and Open-Source Solutions for Automated Image Orientation – A Critical Overview', in: *Progress in Cultural Heritage Preservation: 4th International Conference, EuroMed 2012, Limassol, Cyprus*. pp. 40-54.
- Reu, J.D., Plets, G., Verhoeven, G., Smedt, P.D., Bats, M., Cherretté, B., Maeyer, W.D., Deconynck, J., Herremans, D. and Laloo, P. (2013) Towards a three-dimensional cost-effective registration of the archaeological heritage. *Journal of Archaeological Science*, 40(2), pp.1108–1121.
- Vedaldi, A., Fulkerson, B. (2010) VLFeat - An open and portable library of computer vision algorithms. In: *18th Annual ACM Intern. Conf. on Multimedia*. pp. 1469-1472.
- Wu, C., Agarwal, S., Curless, B. and Seitz, S.M. (2011) Multicore Bundle Adjustment. In: *IEEE Conference on CVPR*. pp. 3057-3064.