From Point Cloud to Textured Model, the Zamani Laser Scanning Pipeline in Heritage Documentation

Heinz Rüther, Christoph Held, Roshan Bhurtha, Ralph Schroeder, Stephen Wessels
Zamani Research Group, Geomatics Division, University of Cape Town,
Heinz.Ruther@uct.ac.za

Abstract

The paper describes the stages of the laser scanning pipeline from data acquisition to the final 3D computer model based on experiences gained during the ongoing creation of data for the African Cultural Heritage Sites and Landscapes database. The various processes are briefly discussed and challenges are highlighted which need to be addressed to develop the full potential of laser scanning. Experiences with fieldwork, scan registration, hole-filling, data cleaning, modelling and texturing are reported. The potential strengths and weaknesses of the emerging tool of "Structure from Motion" are briefly explored for their potential use in combination with laser scanning.

Keyword: Terrestrial Laser Scanning, Cultural Heritage, Digital Preservation, Surface Reconstruction, Texturing, Structure from Motion, Registration, Field Procedures

1. Introduction

The creation of 3D models of architectural structures for monitoring, conservation and restoration interventions at heritage sites poses special challenges for data acquisition and processing. It is not sufficient to represent buildings by a combination of geometric primitives, such as planes for walls, cylinders for columns or spheres for dome shaped roofs. Instead, surfaces need be recorded in great detail in the form of high resolution point clouds. Point clouds can best be acquired photogrammetrically or by laser scanning. Although terrestrial laser scanning has in recent years partially replaced photogrammetry as the principal documentation tool for heritage projects, this technology has not reached a level of perfection which justifies the exclusion of photogrammetry altogether. While there is no doubt that this powerful technology is ideally suited to the task, it is not without its challenges, which arise during the various stages of the pipeline leading from data acquisition stage to the final 3D model. The move from 3D vector models to meshed models based on dense points has received mixed reaction from the wide range of potential end users of the final products. These reactions vary from scepticism to unrealistically high expectations, often arising from unfamiliarity with this new approach or lack of experience with the software for manipulating the generated computer models. High resolution point-cloud based 3D models are generally well received when displayed on the computer screen, however, when it comes to the practical applications of such models, an element of bewilderment often replaces the original enthusiasm. It is therefore important to understand the capabilities and limitations of the technology and relate these to the needs of the 3D data user community. It is crucial that the communication between data producer and the data user is not one sided; it must not only be driven by 'what-can-be-done" but also by "what-is-needed".

The paper attempts to outline the laser scanning pipeline and briefly discusses recent developments in fully automated photogrammetric solution based on "structure from motion" algorithms.

The observations reported in this paper are based on the experiences gained during the modelling of more than 100 historical buildings, archaeological sites and rock shelters produced in the course of the ongoing development of the "African Cultural Heritage Sites and Landscapes Database", a project funded by the Andrew W. Mellon Foundation and developed and populated by the Zamani team based at the University of Cape Town. The paper ends with a brief description of the African Heritage Database.

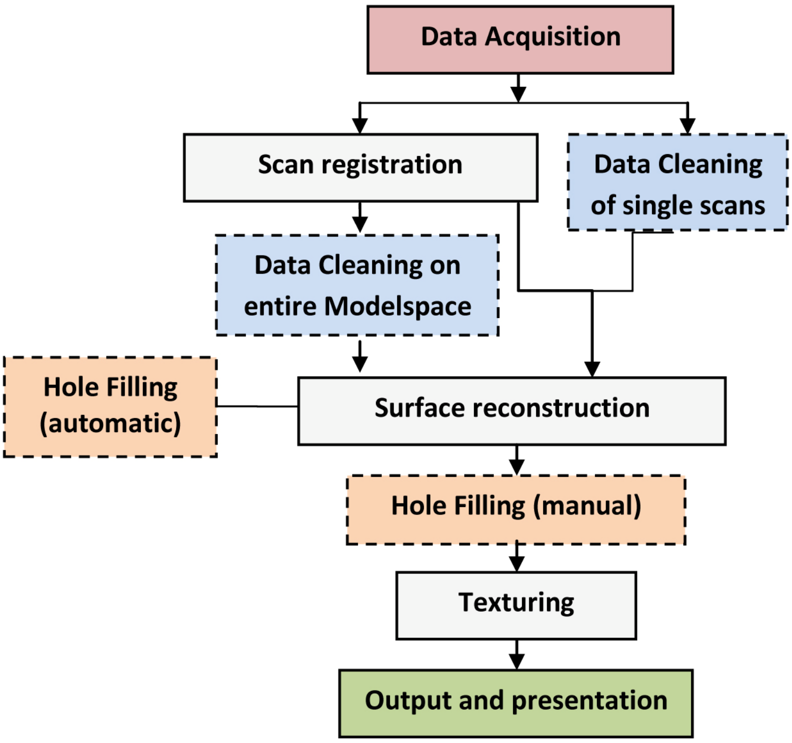

Figure 1. The Laser Scanning Pipeline, optional steps are marked with dotted lines

2. The Laser Scanning Pipeline

The laser scanning pipeline (Fig. 1), the process of creating a 3D computer model from laserscan data, can be subdivided into the stages of data acquisition, registration, data cleaning, modelling, hole filling, texturing and output. Each step of the pipeline will be briefly discussed in the following text and some of the challenges associated with each step will be highlighted.

2.1 Data Acquisition

The development of terrestrial laser scanning from primarily time-of-flight (ToF) instruments to phase-based or combined technology has had an impact on the field procedures in heritage documentation. First field campaigns executed by the Zamani team some five years ago with a ToF scanner produced five to ten full-dome medium resolution scans (2 cm at 20 m distance) per field day. An average site, typically comprising of one major building such as a mosque, church or fortress, was recorded by 40 to 50 scans. Today's field campaigns with phase based scanners can result in more than 1000 individual scans per site, captured at a rate of 80 to 120 per day (Fig 2). Point clouds have increased from between 20 to 50 million points per site to up to 7 billion points. The dramatic increase in the number of set ups over time was a consequence of the development from early time-of-flight scanners with scan times of two to three hours for a full-dome scan to phase based scanners with scan times of three to six minutes for the same resolution. The highest number of scans with this resolution acquired by a single operator of the Zamani Project in one day with a Leica HDS 6100 scanner was 140. This was possible not only because of the high scan rate of the phase based scanner, but also because of the one-button operation and the built in data storage and batteries, which did away with the cumbersome transport of laptop, batteries and cables from station to station, as required for earlier instruments.

Essential for a scanning field team is not only field experience, but, more importantly, experience in all aspects of data processing. Only a clear understanding of the data pipeline and the complexities of each step will guarantee that the principal criteria for a successful field campaign can be met. These include:

- complete point cloud coverage of the structure/site, or at least as complete as physically possible. One has to accept that complete coverage of a monument or site is in most cases unachievable
- appropriate choice of resolution depending on the complexity of the surface, the level of detail required and the variation in distance and angle to the scanned surfaces, bearing in mind that resolution changes with distance from the scanner. For the African Heritage project, buildings are typically scanned with 1- 2 cm point spacing (or higher where need be), while terrain is captured with point intervals varying from 10 to 50 cm
- sufficient overlap for Iterative Closest Point (ICP) based registration, where overlap areas must be chosen to contain sufficient surface detail to allow registration algorithms to find a unique solution. Overlaps should also be chosen to allow for similar angles of incidence from the scan stations

- choice of scanner positions which avoid acute angles of incident for all surfaces, but especially for overlap areas
- choice of an economic target distribution, where targets are used for geo-referencing and quality control. Targets need be visible from a maximum number of instrument setups, without however, sacrificing geometric requirements for stable 3D transformations

The Zamani project has adopted a field procedure in which the scanner is always levelled, and oriented in the same direction. The latter is arbitrarily chosen for the first scan and then maintained throughout the site documentation. The uniform alignment of the scans, while not essential, significantly speeds up the registration process. It proved quite sufficient to maintain consistent scan orientation by simply estimating the orientation, something which can easily be done outside buildings via remote objects and less easily inside. In a number of scan projects the team used RTK GPS equipment to also obtain scanner positions and elevations in open areas. However, the benefit of this additional information for the registration process was marginal when compared to the effort associated with a GPS survey.

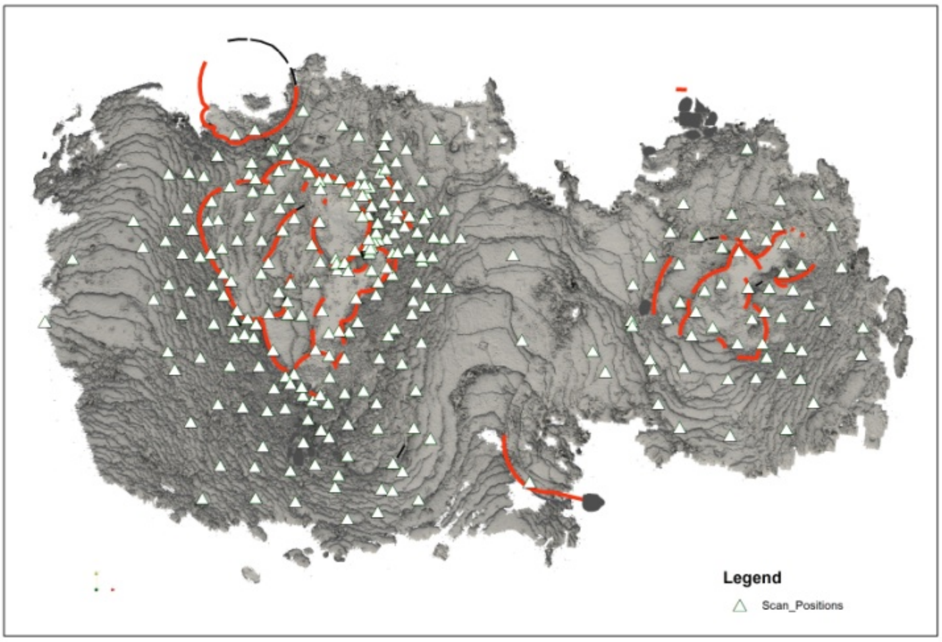

Figure 2. Scanner positions on the terraces of sites DGB1 & DGB 2, Cameroon

2.2 Scan Registration

The first step in producing a 3D model is the registration of the individual scans acquired in the field. This involves the transformation of all scans into a single uniform coordinate system. Registration is a vital element in the data processing pipeline, as the accuracy achieved in this stage influences all subsequent processing steps. There are two approaches to combining the individual scans into a single co-ordinate system. One option relies on targets common to some, or all, of the scans while the other approach aligns scans based on overlapping surfaces, using well established variants of the Iterated Closest Point (ICP) algorithm (Besl and McKay, 1992; Chen and Medioni, 1992; Gruen and Akca, 2005).

Targets are generally easy to identify and locate in point clouds during the registration process, thus reducing processing time especially with software allowing for automatic target detection. Targets also provide high registration accuracies and reduce the danger of mis-registration. Their disadvantage is the requirement of high resolution sub-scans of the area around the target area which in turn increases scan times in the field, unless the entire scan is executed at high resolution. In heritage documentation, however, one encounters a number of practical problems. It is often difficult, if not impossible, to physically place targets on fragile or high walls and in difficult-to-access areas. It is possible, however, to locate targets outside the scanned object, provided that they are visible from more than one scan position. But this option is obviously not available when scanning the inside of buildings. A further argument against targets is the complexity of heritage buildings. For example, the so-called 'palace' at the ruined Swahili town of Songo Mnara in Tanzania comprises of more than 100 rooms. Considering the need for a minimum of three and preferably more targets for each scan, the use of targets in such environments is highly impractical.

The policy adopted for the scanning of sites for the African Heritage documentation project is therefore to rely on surface based registration and to use only limited numbers of full-sphere targets for very large sites which are placed in exposed positions where they can be viewed from a maximum number of scans. The coordinates of these targets are determined by post-processed GPS surveys wherever possible and they serve to check and position the final model. For some sites point clouds were registered with both, target and surface based methods, but no significant difference between the accuracies of the resulting registered point clouds could be detected.

Soft- and hardware limitations become obvious when registering large numbers of scans, which require long processing times and can cause system crashes. The Zamani team first creates a skeleton model of the entire site or structure by registering a minimum number of individual, usually, longer distance scans. This is followed by a global registration of the skeleton model, provided the data volume allows this. The skeleton model is then filled in with the bulk of short range scans.

2.3 Cleaning of Point Clouds

The cleaning of the data, i.e. the removal of objects not relevant to the documentation, is an important step in the 3D modelling pipeline. As a rule the Zamani team cleans individual scans of unwanted objects such as trees, people, cables, cars, doors, animals and random objects prior to the final modelling process. Pre-model cleaning results in smaller models, especially on sites with dense vegetation (Fig. 3). Only minor remaining flaws are dealt with on the model.

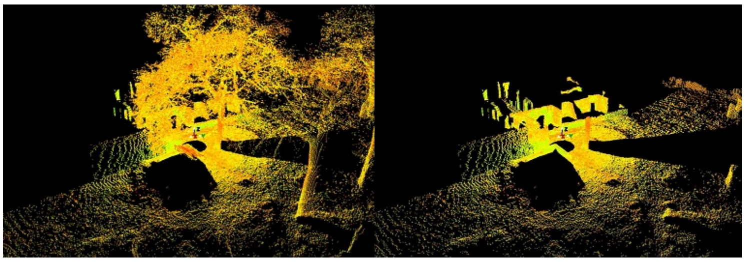

Figure 3. Uncleaned point cloud with 6.7 million points (left), cleaned point cloud with 4 million points (right)

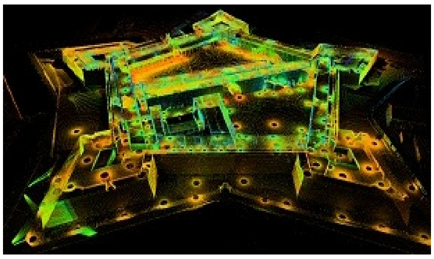

Figure 4. Screenshot of the point cloud of the Castle, Cape Town, South Africa

One can also clean the fully registered point cloud (Fig. 4). This reduces the cleaning time, as objects appearing in multiple scans can be removed in one operation. However, problems arise in the subsequent meshing stage as the grid structure of individual scan files is lost in overlapping areas after all points are merged into a single cloud. The huge data volume of a combined point cloud also results in frequent system crashes or in long waiting times when manipulating the point cloud during cleaning on the computer screen.

While some artefacts in scans can be removed automatically most cleaning of scans is still a manual, tedious and time consuming task. A typical scanning project in the African heritage project comprises of 300 – 800 scans each of which typically requires 0.5 to 2 hours of cleaning time by an experienced operator, while complex and very high resolution scenes can take a full day of cleaning per scan. Here, further development of automatic or semi-automatic cleaning tools, capable of working with large datasets could greatly improve the workflow.

2.4 Surface Reconstruction

Meshing, the conversion of the discrete point set to a continuous surface in the form of triangles, is the next step in the pipeline. Surfaces can either be formed by directly connecting points which then become part of the surface (Amenta, Choi and Kolluri, 2001; Bernardini et al, 1999) or by creating a best-fit surface through the points (Curless and Levoy, 1996; Hoppe et al, 1992; Kazhdan, Bolito and Hoppe, 2006).

The former approach usually employs the Voronoi diagram and the corresponding Delaunay triangulation to find point neighbours and create connections. This method allows the creation of the model from a complete, unstructured point cloud, but it has the disadvantage of being very susceptible to instrument noise, outliers and scan alignment errors. Noise turns a surface, which in theory should be infinitely thin, into a ‘sheet’ of finite thickness. Connecting the closest neighbours in this 'sheet' without filtering can result in incorrect triangles and faces which are non-consistent in their orientation to their neighbours. Thus a noise-filtering pre-step becomes necessary. Depending on the dataset, this can be a lengthy and complicated process. Noise reduction is generally indiscriminate and unfortunately accompanied by a loss of detail. Noise typically increases in areas, where scans overlap or as a result of small errors in registration, which create multiple surfaces.

Best-fit surface generating algorithms automatically perform noise reduction by approximating the surface and thus remove noise. Other approaches employ pre-triangulated scans to extract a surface by averaging between overlapping areas and thus smoothing the input data. The danger in all smoothing operations is the potential loss of relevant detail, which might be problematic when inspecting the condition of surfaces or monitoring erosion phenomena.

Depending on the source data, calculating vertex normals can be very trivial or very complex. It is trivial when the data is provided in grid form, which is the case for individual scans. In this case the normals can be retrieved by using the scanning grid information which indicates nearest neighbours. If there is no such information available, or if two or more overlapping scans are aligned and merged into one cloud, the grid structure is lost and nearest neighbours have to be found in time consuming computing processes.

In general, a compromise has to be found between reliability and completeness of the model, as 'completeness' might require the inclusion of uncertain data, which could, for example, be created by an inappropriately set threshold, which decides whether a point's neighbour is close enough to be used. Setting this threshold too large will create artefacts and connect objects which should not be connected, or, if set too small, will create holes, even though the surface was sufficiently covered by scanned points.

This is especially problematic for individual scans required by some volumetric approaches with greatly varying object distances from the scanner. As sampling density decreases with distance from the scanner, point connections and normals become increasingly unreliable. This problem is reduced when working with complete point clouds, thus considering points from other scans. In this way, sampling density is increased in previously sparse areas resulting in a more complete and still reliable model.

It would seem that none of the currently available modelling algorithms produce a perfect model. At present, a combination of multiple versions of the same model, created with different algorithms, e.g. Poisson (Kazhdan, Bolito and Hoppe, 2006) and Volumetric Integration(Curless and Levoy, 1996), appears to lead to the most complete and detailed model. The process of merging different models is not trivial as the various models can differ significantly in resolution and detail (Fig. 5a, 5b).

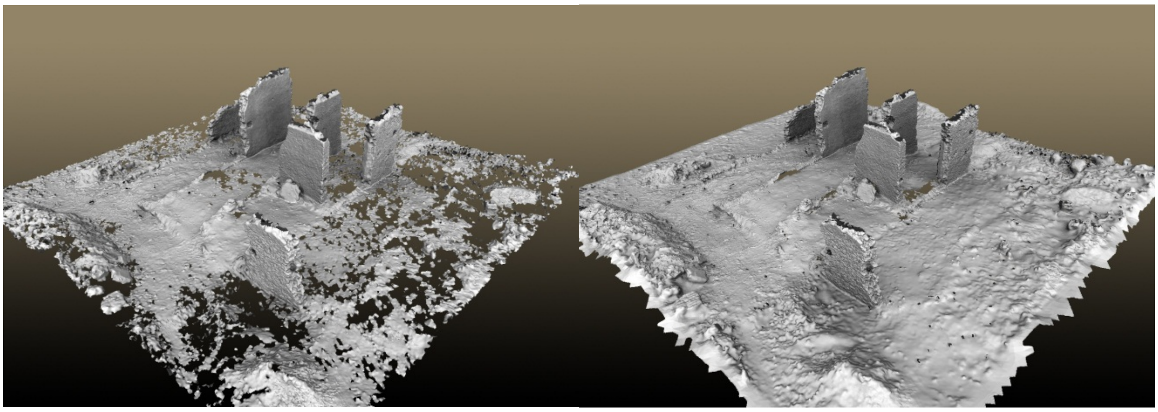

Figure 5a (left). Model meshed only with the volumetric integration method, Figure 5b (right). Combination of the Poisson algorithm for the ground and volumetric integration of the walls.

A further possible cause for artefacts in the modelling process is the nature of the scanned surface. For example grass, leaves and branches of bushes or trees or moving objects like curtains, cars and people should be cleaned out before creating the model. While the latter are generally easy to identify and remove, grass and leaves are more complex, since they are typically only represented by a few points which makes identification difficult. They might be modelled by the meshing algorithm to produce only a single sided surface with a complex boundary, since a blade of grass or a leaf of a tree or bush is hardly ever scanned from two sides. This is difficult to repair at a later stage. Current vertex-normal-based implementations (Fiorin et al, 2007; Bolitho et al, 2007), tested by the project, have a good and reasonable smoothing effect on grass without smoothing out too much important detail on the objects of interest, but unfortunately they are not very robust when applied to such irregular surfaces and often crash.

Meshing the large models encountered in the African Heritage project cannot be done with standard, in-core software and thus out-of core tools, and /or streaming techniques with the automatic ability to split up data, are becoming essential. The present approach of the Zamani team is to split the data into subsets to be used for parallel processes with out-of-core tools. Overlaps at subset boundaries are introduced to allow smooth connections of neighbouring sets.

This way, the models for the African Heritage project are usually extracted at a resolution of 2 cm (Callieri et al, 2003) with a software variant of the Volumetric Integration approach, but much higher resolution models from the raw data can be generated, if the need arises (Fig. 6).

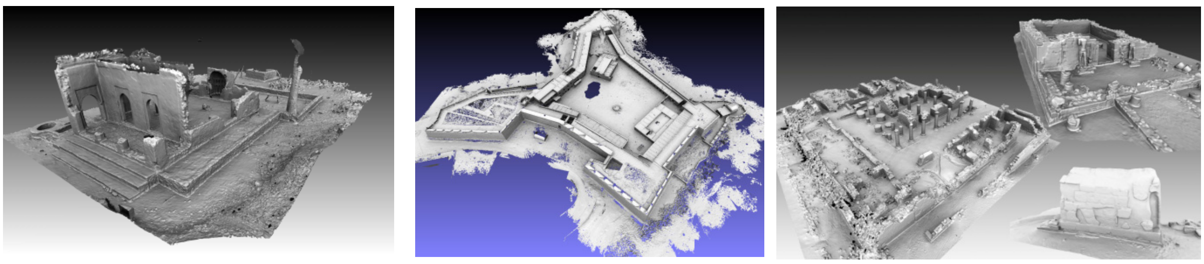

Figure 6. 3D model screenshot of a mosque on Songo Mnara, Tanzania (left), 3D Model screenshot of San Sebastian Fortress on Mozambique Island, Mozambique (centre), 3D model screenshot of Musawwarat es- Sufra, Sudan (right)

2.5 Surface Augmentation

Only in very exceptional cases, if ever, is a laser scan model free of scan holes. Holes occur especially when scanning very complex and irregular objects, or wherever a surface is invisible to the scanner. Typical examples are ornamental building facades or upward facing surfaces, where no scan positions can be found above the surface, such as window ledges and roofs. The same problem arises when acquiring photography for texturing the model.

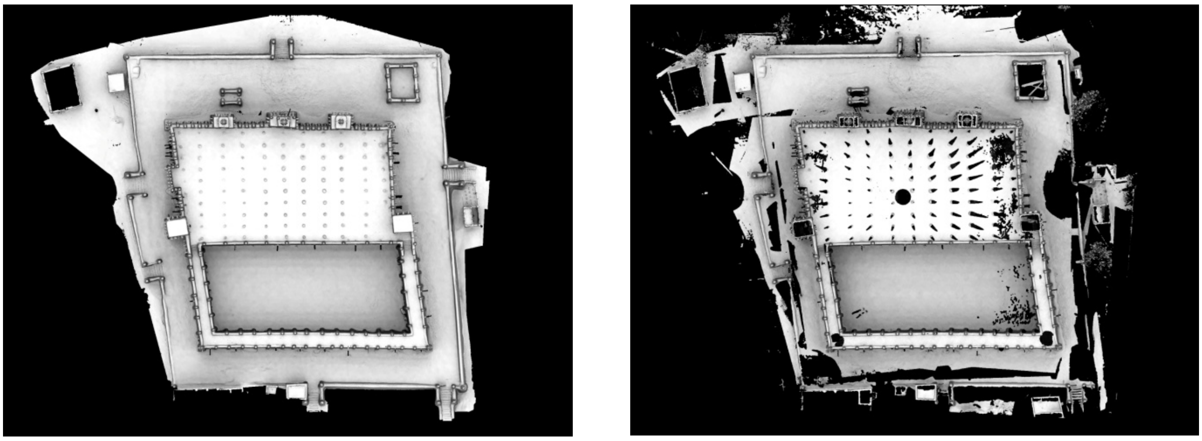

Figure 7. Filled 3D model of the Djenne Mosque, Mali (left). Unfilled 3D model of the Djenne Mosque (right).

If no scan data are available, current software fills in small holes automatically and offers conventional modelling or cloning options for larger patches (Fig 7). This semi-automatic approach is a time intensive task, since it is heavily user-based.

Methods for automated hole-filling or surface augmentation have been developed (Sharf, Alexa and Cohen-Or, 2004) but their use in heritage documentation is questionable. Surface augmentation algorithms can fill holes plausibly, which makes it difficult and even impossible to distinguish between real surface data and artificially introduced patches. This is not acceptable for the scientific documentation of cultural heritage sites, where data might be used for research or restoration. On the other hand, hole free surfaces are aesthetically more appealing and necessary to produce interactive 3D walkthroughs. Watertight models are also required when producing physical to-scale models with a 3D printer. It would seem desirable for the heritage scanning community if software, performing hole-filling and model-viewing, could make use of a standard display format clearly indicating augmented surface portions, on request.

Scan holes in roof areas, which are a major challenge in the documentation of buildings, can be closed photogrammetrically through images taken from Unmanned Aerial Vehicles (UAV's). Unfortunately the Zamani Project does not have access to a UAV.

2.6 Photorealistic Texturing

Colour can enhance a 3D-Model not only visually, but also assist with interpretation and diagnosis. Photorealistic colour of a surface can, for example, assist with the detection and monitoring of eroded, chemically changed or restored surface areas. While textured models might be desirable, but not essential, for many applications, they are critical when modelling rock art shelters (Fig. 8).

Figure 8. Textured 3D model of rock art panel at Game Pass shelter, Drakensberg, South Africa.

Laser-scanner manufacturers increasingly equip their instruments with photo- or video cameras to achieve the colourisation of the scanned surface. These built-in systems are convenient to use, but so far they do not seem to reach the quality of independent, external cameras. Images tend to have a blurry or milky appearance. Some manufactures thus offer adapters for external cameras which take a panoramic image from the position of the scanner. The biggest disadvantage of these approaches is the inconsistent lighting between scanning positions. A scanning campaign of a heritage sites usually spans several days with the scanner operating throughout the day and in some cases at night. This necessarily leads to scanning under very different light conditions and if the images are taken concurrently with the scans, they can differ significantly throughout the model. This suggests that the optimal way to acquire texturing photography for a model is independent photography, taken over a minimal time period, or on different days with approximately the same lighting conditions.

To place an image onto a 3D-Model, the camera parameters, internal and external, have to be known. An accurate projection can only be accomplished with the correct position and orientation of the camera in space, the correct focal length and distortion free images. The best approach is to use images acquired with a calibrated camera and lens. But this method also has its disadvantage as it requires a preset fixed focus, which might result in some of the captured images being blurred. Theoretically the camera can be calibrated individually for each photo but this is extremely time-consuming, whatever technique one uses for the calibration.

The internal as well as the external parameters can also be estimated or post-calibrated. If the parameters are not known at all, current software asks the user to find corresponding points on the model and on the image. But even if the amount of correlated points is reduced to a bare minimum, these points need to be chosen accurately and thus, covering a large structure completely in detail with hundreds or even thousands of images, will take a significant amount of time. Full-dome panoramic images, converted to cube maps, can be used to reduce the amount of images significantly. Per panorama, only one sub-image needs to be registered, while the remaining images differ only in their orientation. Panoramic images are nearly free of lens distortions, since the camera parameters are estimated when stitching the panorama image. Computational tools to refine initial camera parameter estimations and thus increase the accuracy of the image-model-fit have been successfully demonstrated (Corsini et al, 2009; Gal et al, 2010).

Current texturing techniques encourage the user to minimise the number of photos to reduce processing time. However, advances in structure-from-motion algorithms (Snavely, Seitz and Szeliski, 2006) suggest the acquisition and use of large numbers of images. These algorithms require large overlaps and small separations between camera positions and can almost automatically estimate all camera parameters required for texture mapping. Tests within the project have shown that using SfM for texturing is an effective and fast way of aligning photos, covering a large portion of the objects automatically. However, blending a large number of images can result in a blurred texture, especially when using a non-calibrated lens.

If the task of finding the camera parameters is solved, it has to be decided how to project the images onto the surface. There are two ways, either by projective texture mapping or by assigning colours to the vertices. The first approach requires that each vertex (3D model coordinates) on the model is associated with a corresponding point (2D image coordinates) on one or more of the images. It is not trivial to handle these extremely large datasets of extensive models with hundreds of associated images.

It is much easier to assign a colour directly to each vertex, and then interpolate the colour on the surface between the vertices. This requires that the original point cloud resolution is high enough to represent the required detail. If this is not the case one can artificially densify the point cloud by subdividing the original triangles. In this case each new point needs to be located in the image set to obtain its corresponding 'real' colour. Practically, this requires tools which allow the loading of large models and to then further increase their size. At present such tools appear elusive.

3. Structure From Motion Methods

The recent emergence of fully automated photogrammetric software (Bundler/ PMVS2, AgiSoftStereoScan/AgiSoftPhotoScan or Photosynth) based on Structure from Motion (SfM) algorithms makes it possible to create textured 3D point clouds or even meshed models from un-calibrated, unordered images.

Figure 9: 3D model using SfM of a Kasbah in Morocco

This method is developed by the computer graphics community, based on photogrammetric and computer vision algorithms. It allows the user to enter into the software application a collection of images captured with one or more handheld, un-calibrated camera(s) without any information about the relative position or orientation of the imagery or their relation to each other. The software then fully automatically associates the images to each other, determines their position and orientation in space and creates a point cloud. Some software, such as bundler (Snavely, Seitz and Szeliski, 2006) only creates sparse point clouds which are enriched with dense points by a second application (PMVS2). Other software (e.g. AgiSoft Photoscan) create a dense mesh in one process. The vertices of the point cloud/mesh are automatically associated with RGB values taken from the photography. The point cloud is un-scaled and external information must be called upon to scale the model.

The only skill required in the process is the choice of positions for the camera while capturing the imagery. A simple rule is that images should be taken sequentially at short intervals while moving around the object. A structured sequence is not essential but speeds up the search time during the image association stage of the process.

The African heritage project experimented with SfM applications with mixed results. Results are unpredictable and vary from excellent (Fig. 9) to poor and metric quality, reliability and robustness are not guaranteed.

The role of SfM in the African Heritage will at first be restricted to fill small surface areas not reachable by the scanner. As the results of the SfM algorithm are dense point clouds or even triangular meshes it is possible to register the acquired SfM models to the laserscan model using standard registration software, thus providing scale and quality control. SfM in its present form appears unsuitable to the documentation of complex buildings with large numbers of rooms such as the palaces and fortresses mentioned above.

There are ongoing projects, based on traditional photogrammetry which can be expected to provide the same automation and functionality as SfM but with the statistics- based accuracy and reliability analysis inherent to traditional photogrammetry. This development clearly deserves attention (Barazzetti, Scaioni and Remondino, 2010).

4. The Final Product and the End-User

A discussion between photogrammetric and heritage professionals, dating back to the first photogrammetric documentation of heritage sites and specifically rock art, was focussed on the use of photographic images as opposed to hand drawing and tracing. This discussion seems to have re-emerged between the creators of 3D computer models on the one side and some architects and conservators on the other. There appears to be a condescending attitude to digital technologies which are sometimes seen as sterile and void of the possibility of interpretation. From an objective perspective such absence of subjective interpretation would appear to be an advantage for scientific analysis rather than a shortcoming. The irony of this discussion is that computer generated models are also not entirely objective, with final results depending to a not insignificant extend, on the modelling software used. It is important to make end-users aware that it is impossible to create absolutely correct models, even if point accuracies and resolutions are high, interpolation, smoothing and artefacts are unavoidable and end-results will not be without flaws. Nevertheless, it can safely be said that the level of objectivity achieved by digital documentation, if applied in a scientifically rigorous way, is higher than that of manual techniques.

5. The African Cultural Heritage Sites and Landscapes Database

The African Cultural Heritage Sites and Landscapes Database project is dedicated to recording, and thus contributing to the protecting of, cultural heritage monuments and landscapes throughout the African continent. The project recognizes the urgent need to:

- create metrically correct digital (3D) documentations of African Heritage Sites for future generations
- provide spatial information of heritage sites for conservation and restoration
- provide spatial information of heritage sites for education and research
- promote awareness of Architectural African Heritage (historical sites, buildings and towns) within Africa and worldwide
- provide management tools for site management at local and regional level

The project was initiated with a grant from the Andrew W. Mellon foundation and is executed by the Zamani Research Group, based at the University of Cape Town (UCT). The project uses state-of-the-art technologies, such as Laser Scanning, GIS, close-range and aerial Photogrammetry, Remote Sensing, as well as virtual reality technology, to create an integrated database of important African cultural heritage sites. The output of the project is made available by the JStor digital library, New York (http://JStor.org and http://aluka.org). Full data sets are available for research, education, conservation and restoration projects. A small subset of the data can be viewed on http://www.zamaniproject.org.

The Zamani group has to date documented some 45 African heritage sites in Ghana, Mali, Mozambique, Tanzania, Kenya, Ethiopia, Sudan, Egypt, Algeria, Zimbabwe, Cameroon, Uganda and South Africa. All projects are executed by members of the unit, with the support of staff members of Antiquities or equivalent Government Departments. A total of approximately fifteen terabytes of data on African sites has been generated by the Zamani group over the past six years. These data are augmented by relevant contextual information data selected by JStor/Aluka.

All projects are carried out in close cooperation with local museum and heritage authorities as well as staff and students of local universities and the University of Cape Town. The generated information and database has been disseminated through seven workshops, the JStor/Aluka website and a much smaller UCT website.

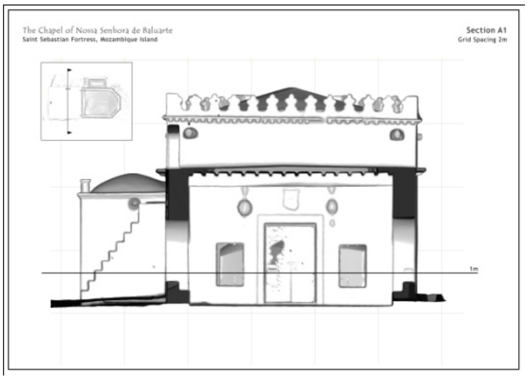

Figure 10: Section of Chapel of Nossa Senhora, Mozambique

Figure 11: 3D model of the Silk Tombs in Petra, Jordan

The deliverables of the project are:

- an integrated database consisting of a Spatial/Geographic Information System (GIS) for each of the sites
- 3D computer models of structures and parts of towns and landscapes (Fig. 11)
- elevations, sections, ground plans and roof plans (Fig. 10)
- computer visualisations with walk-through and other inspection capabilities (where feasible)
- spherical panorama tours
- contextual photographs
- site related digitized documents, scientific papers, excavation reports and similar material

During the process of creating the database, a methodology for the documentation of African heritage sites has been developed and optimal ways are explored in which the data can be used by African heritage authorities and by museum officials and researchers in Africa and worldwide.

6. Conclusion

Laser scanning has established a prominent position in the spatial documentation of heritage sites and rightfully so. It is a powerful tool which makes it possible to collect valuable and accurate spatial information in relatively short field time. It played a prominent role in the acquisition of data for the African Heritage Database and it is unlikely that other techniques would have been able to acquire similar data volumes. However, during the course of the African Heritage Project, it also became obvious that laser scanning alone does not hold all the answers to heritage documentation. Photogrammetry still has an important role to play, both, using cameras directly linked to the scanner or using traditional close-range photogrammetry, with independent cameras. Both approaches can contribute to texturing and feature extraction, while close-range photogrammetry is also essential for the capturing of surfaces, which are difficult to access by scanners. It is also obvious that there is still room for development in the many complex steps in the laser scanning pipeline, from data collection and scanner specifications, registration, cleaning, meshing and hole-filling, and texturing.

Visually appealing textured 3D computer models, video fly-throughs and interactive presentations appear attractive to the end user and may be preferred to incomplete models with remaining scan holes. Unrealistic expectations with respect to 3D model quality, resolution, size and texturing have been raised through internet publications of seemingly perfect models, thus making it difficult for projects to meet user expectations. Because of this, unintended pressure can be placed on the producer of such data to sacrifice authenticity and accuracy in favour of visual appeal by, for example, filling holes, creating unrealistic lighting effects and hiding low resolution by texture. This is contrary to the ethos of Heritage Documentation, which requires objective, accurate and unmodified data. It would therefore seem imperative to develop specifications which clearly define strengths and limitations of a generated model which can be conveyed to the user. Without such descriptors, the scientific, engineering and cultural worth of these models may be reduced. However, technical specifications are only then meaningful, if the users are aware of the endemic limitations of this particular methodology. Furthermore, it is important to explore new ways of using laser scanned models, besides for visual presentation, to justify the effort and cost spent on their creation.

7. Acknowledgements

This African Heritage database project has been made possible with the generous support of the Mellon Foundation, the assistance of Rahim Rajan from Aluka/JStor, and the invaluable advice of the team of experts on African history – Dr. George Abungu, Prof. Susan Macintosh and Prof. Martin Hall. The authors thank iQlaser, Leica Geosystems, and Optron Geomatics, who have generously provided additional laser scanning equipment for field campaigns. Our gratitude goes to the Visual Computing Lab of the ISTI-CNR, Pisa, Italy, Stanford Computer Graphics Laboratory, USA and to InnovMetric Software Inc, Canada. A final thank you goes to the University of Cape Town for housing and supporting the project.

8. References

Amenta, N., Choi, S. and Kolluri, R. K. 2001, 'The power crust', in K Lee and D Anderson (eds.), *Proceedings of the 6th ACM Symposium on Solid Modeling and Applications*, Ann Arbor, 04 – 08 June 2001, pp. 249–260, ACM, New York, 2001.

Barazzetti, L., Scaioni, M. and Remondino, F. 2010, 'Orientation and 3D modelling from markerless terrestrial images: combining accuracy with automation', *The Photogrammetric Record,* vol. 25, no. 132, pp. 356--381.

Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C. and Taubin, G. 1999, 'The ball-pivoting algorithm for surface reconstruction'. *IEEE Transactions on Visualization and Computer Graphics, vol.* 5, no. 4, pp. 349–359.

Besl, P.J. and McKay, M.D. 1992. A method for registration of 3D shapes*, IEEE Trans. Pattern. Anal. and Mach. Intell*, vol. 14, no 2, pp. 239-256.

Bolitho, M., Kazhdan, M., Burns, R. and Hoppe, H. 2007. 'Multilevel streaming for out-of-core surface reconstruction', in A Belyaev and M Garland (eds.), *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing*, Barcelona, 04-06 July 2007, pp. 69–78, Eurographics Association, Aire-la-Ville, 2007.

Callieri, M., Cignoni, P., Ganovelli, F., Montani, C., Pingi, P. and Scopigno, R. 2003. 'Vclab's tools for 3d range data processing', in D Arnold, A Chalmers and F Niccolucci (eds.), *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST2003) and First EUROGRAPHICS Workshop on Graphics and Cultural Heritage*, Brighton, 05-07 November 2003, pp. 13-22, Eurographics Association, Aire-la-Ville, 2003.

Chen, Y., and Medioni, G. 1992. 'Object modelling by registration of multiple range images', *Image Vision Comput.*, vol. 10, no. 3, pp. 145-155.

Corsini, M., Dellepiane, M., Ponchio, F. and Scopigno, R. 2009. 'Image-to-Geometry Registration: a Mutual Information Method exploiting Illumination-related Geometric Properties', *Computer Graphics Forum*, vol. 28, no. 7, pp. 1755-1764.

Curless, B., and Levoy, M. 1996. 'A volumetric method for building complex models from range images', in J Fuji (ed), *SIGGRAPH 96 23rd International Conference on Computer Graphics and Interactive Technique*, New Orleans, 04-09 August 1996, pp. 303–312, ACM, New York, 1996.

Fiorin, V., Cignoni, P. and Scopigno, R. 2007. 'Out-of-core mls reconstruction', in E Gobetti (ed.), *Proceedings of the Ninth IASTED International Conference on Computer Graphics and Imaging - CGIM 2007,* Innsbruck, 13-15 February 2007, , pp. 27–34, ACTA Press, Anaheim, 2007.

Gal, R., Wexler, Y., Ofek, E., Hoppe, H. and Cohen-Or, D. 2010, 'Seamless Montage for Texturing Models', *Computer Graphics Forum*, vol. 29, no. 2, pp. 479-486.

Gruen, A. & Akca, D. 2005, 'Least squares 3D surface and curve matching', *ISPRS Journal of Photogrammetry and Remote Sensing, vol.* 59, no. 3, pp. 151-174.

Hoppe, H., DeRose, T., Duchampt, T., McDonald, J. and Stuetzle, W. 1992, 'Surface Reconstruction from Unorganized Points', in J Thomas (ed.), *Proceedings of the 19st Annual Conference on Computer Graphics and Interactive Techniques, Chicago, 27-31 July 1992, SIGGRAPH*, pp. 71-78, ACM, New York

Kazhdan, M., Bolitho, M. and Hoppe, H. 2006. 'Poisson surface reconstruction', in A Sheffer and K Polthier (Eds.), *Proceedings of the 4th Eurographics Symposium on Geometry Processing*, Cagliari, 26-28 June 2006, pp. 61–70, Eurographics Association, Aire-la-Ville 2006.

Sharf, A., Alexa, M. and Cohen-Or, D. 2004. 'Context-based surface completion', in J Marks (ED.), *ACM SIGGRAPH 2004 Papers*, Los Angeles, 08-12 August 2004, pp. 878–887, ACM, New York, 2004.

Snavely, N., Seitz, SM. and Szeliski, R. 2006. 'Photo tourism: exploring photo collections in 3d', *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2006*, vol 25, no. 3, pp. 835-846.