

Multiview 3D Reconstruction of the Archaeological Site at Weymouth from Image Series

Benjamin Ducke, David Score, Joseph Reeves
Oxford Archaeology

Abstract

Multiview (n-view, or multiple view) 3D reconstruction is the computationally complex process by which a full 3D model is derived from a series of overlapping images. It is based on research in the field of computer vision which in turn relies on older methods from photogrammetry. This report presents a multiview reconstruction tool chain composed from various freely available, open source components and presents a practical application example in the form of a 3D model of an archaeological site.

Keywords

Computer vision, archaeology, photogrammetry, multiview, 3D reconstruction, open source, Weymouth

1. Introduction

Archaeological sites can exhibit considerable structural complexity. Documenting them accurately and efficiently in three dimensions is a technological challenge that needs to be addressed not only for the sake of producing richer scientific datasets but also because archaeology in the media age needs to provide presentations of its work to the general public that are more attention catching than flat, technical illustrations. Software technology that can create full 3D models from simply the photographic coverage of an object or scene is therefore of great interest to archaeological site documentation and presentation. After all, comprehensive series of images are taken routinely on all excavation projects.

In June and July 2009 during the earthwork operation for the construction of the Weymouth Relief Road in Dorset, archaeologists from Oxford Archaeology, working for Skanska Construction and Dorset County Council, excavated an extraordinary burial site which was discovered on the crest of the Dorset Ridgeway (Fig. 1). 51 decapitated skulls had been placed in a pile in a disused quarry pit approximately 8 m in diameter and their associated bodies had been discarded haphazardly in another area of the same pit. The remains were dated to between AD910 and AD1030 in the Anglo-Saxon period and later isotope analysis showed them to be of Scandinavian origin – Vikings.

This important and complex bone deposit was excavated in detail using a variety of archaeological techniques which included both analogue and digital photography. The digital image archive was not taken with 3D reconstruction in mind, but given the particular nature of the find and the prolonged time-scale for excavation, it was felt that the extensive record which had been made would provide a good opportunity to trial and assess the software currently available for multiview 3D reconstruction.

The work presented here was initiated by Oxford Archaeology to test the feasibility of "Bonus 3D", i.e. adding value and getting the most out of a project's existing digital resources without producing unreasonable extra cost or operational overhead. Terrestrial laser scanners are not universally useful here, as they remain expensive and their operation complex and time-consuming. Similar problems hamper many common photogrammetry solutions with prohibitively high licensing costs and a reliance on time-consuming, manual workflows. As an alternative, we discuss a largely automated processing method that relies on a set of freely available software tools.



Fig. 1: The excavations at Weymouth in progress.

2. Public Archaeology in 3D

Materials and results from the Weymouth Relief Road excavation, in particular the mass grave, were displayed at a public exhibition during March 2010 which attracted over 7000 visitors (Fig. 2). The exhibition provided an excellent opportunity to utilise digital 3D reconstruction as a tool for disseminating the nature of the mass burial to a wider audience, the majority of whom were non-specialists in the field and who received much of their archaeological exposure via television or the Internet. Representations of the mass burial were also very important as it formed the primary “find” of the excavation; small finds or artefacts, the type of objects most often displayed because of their unique or representative importance, were not abundantly recovered from the Weymouth mass grave. In this case, therefore, the mass grave as a complete entity had to be represented. The process of recording the mass grave also led to its inevitable destruction, as such a complete reconstruction was paramount in our efforts to display the archaeology of the site.

The public exhibition called for the 3D reconstruction to be displayed as both a computer model and as a selection of still images for use in presentations or on posters. Furthermore, the concept of “Bonus 3D” could extend beyond the exhibition and we intend to use produced materials for both online and print publication.

Whilst the Weymouth excavation shows that an excavated object, completely exposed at one point in time, can be recorded and displayed using the software and techniques detailed above, it also acts to

demonstrate the potentials of other archaeological uses of the technology. Not only was the Weymouth mass burial pit conveniently open for photography, but its spatial structure also made it conducive for this technology; shaped like a gently curved bowl with skeletal elements in sharp relief, it is very suitable for display in 3D. Most archaeological sites are not uncovered in such an ideal way, but the CV techniques used at Weymouth are potentially applicable to all kinds of shapes and objects, such as those within the confines of an excavation trench, or larger objects of interest within a landscape.

In fact, an archaeological deposit may have remained covered since antiquity and exposed only in small sections during the archaeological process; the 3D CV methods described here still give us the ability to photograph it throughout the excavation period and to ultimately produce a model of a deposit or object not seen in its entirety since it was originally buried. Taking the technique to the other extreme, the opportunity exists to create 3D models for the comprehensive presentation of individual finds and artefacts.



Fig. 2: Public display of the Weymouth mass burial.

3. Technical Details

Photogrammetry is the art of determining the geometric properties of objects that are visible on photos. As a collection of mathematical methods, it is as old as photography itself and has for a long time provided tools to produce stereo views, elevation models and rectified aerial images (Mikhail, Bethel and McGlone 2001). It also has a long history in archaeological field work, where it has been used to produce undistorted views of building façades, excavated cross sections and other objects of interest.

From a computer science point of view, CV is the modern field of research that supersedes photogrammetry, taking its geometric foundations and supplementing them with automated, computationally intense procedures (Hartley and Zisserman 2004). Classic photogrammetric tasks, such as the production of orthophotos, represent only a small subset of CV applications which include automated image matching and classification, object tracking, navigation and, as discussed here, full 3D reconstruction from images.

CV plays an integral role in the design of autonomous devices and vehicles and consequently focuses on automated and robust methods, relying on massive input sizes, statistical relationships and raw computational power to compensate for the absence of human judgement and decision making. For example, whereas in a typical photogrammetric workflow a human operator carefully selects a set of control points on a small number of images to establish good correspondences, a CV program will extract thousands of such points from a large number of images, while automatically establishing correspondences between them. While individual control points may be of lesser quality, their combined information value will be superior. So much superior, in fact, that problems such as camera calibration and lens distortion correction can be solved "on-the-fly", at no extra cost.

3.1 An Open Source Approach

Our work is certainly not the first to employ image-based 3D modelling applications in archaeology and cultural heritage management (see Anderson 2010, Campana and Remondino 2008, El-Hakim et al. 2008 for some recent examples). But looking at the limited scale and speed of adoption of these technologies in our field of work, it now seems clear that there will be little sustainable progress and no broad impact in actual practice, unless the algorithms and software implementations are freely available, so that others can put published research into practice, and reproduce, modify, improve and freely disseminate the results.

It therefore does not suffice any longer to publish images and descriptions of software. Rather, it is necessary to publish the software itself, in the form of liberally licensed open source code. Our work may be regarded as an attempt to fully implement these ideas in our specific field of research. All software used in this study is freely available under an open source license.

3.2 General Workflow

The process of multiview 3D reconstruction can be broken down into a number of smaller tasks. CV provides a diversity of algorithms for each of them:

1. Extraction of features (keypoints) from the input images.
2. Image matching and camera reconstruction (Structure from Motion).
3. Dense 3D model reconstruction.
4. Surface reconstruction ("meshing"), manual cleaning and model publication.

In step 1, characteristic regions (features) are automatically identified on all input images. Next (2), these features are tracked across the images and their motion is analyzed to extract the scene's basic 3D geometry. The result will be a sparse 3D point cloud which is then densified (3) by reconstructing a great amount of additional points. Finally (4), the points need to be connected into triangulated

approximations of the original object surfaces. For the purposes of the Weymouth exhibition a fifth stage was added, the rendering of an animation from the model.

Most of the workflow can be fully automated and will still provide good results in almost all cases. Only the final, surface reconstruction step, requires some interactive model editing. In the following sections, we will give some more insight into each of these tasks and provide links to freely available, open source software that can be used to solve them.

3.1 Feature Extraction

Feature extraction is the process that automatically detects characteristic “regions of interest” in an image. To simplify the processing, it is frequently carried out on a greyscale version of the original input image. In that case, regions of interest are those where notable changes of grey values can be observed within a relatively small area. They are defined by location (in pixel coordinates), size and direction of the greyscale gradient. Perhaps the most universal feature extractor is the Scale Invariant Feature Transform (SIFT) algorithm (Lowe 2004). Matching corresponding features in different images with each other is the first step in 3D reconstruction and features extracted with SIFT provide a robust base for this. Typically, SIFT will detect tens of thousands of features even in a relatively low-resolution image and hundreds of thousands in a 10-15 megapixel image.

The first implementation of SIFT was written by its inventor, David Lowe, but not released in open source form. This in itself is not a problem, as there are many alternative open source implementations. However, free use of SIFT is hampered (at least in North America) by US Patent 6,711,293, assigned to The University of British Columbia. The latter seems to be willing to allow royalty free use of its patent for educational and research use. But archaeological contractors and heritage agencies active in North America will need to clarify their eligibility before using SIFT in their work.




Fig. 3: Grayscale image (left) and location, direction and extent of features detected by SIFT (right). Many hundreds of features were detected on this low-resolution image.

3.2 Structure from Motion

Objects in the background move more slowly than those in the foreground. The human brain exploits this simple principle when it generates depth information. The same basic idea is at the heart of CV's Structure from Motion (SFM) approach. Once corresponding features have been identified across a

series of images, their relative speed and direction of movement across the image series provide clues about their position in 3D space. In combination with the focal length and CCD width of the camera used to take the image (to provide information for pixel scaling), this allows for precise reconstruction of the original camera position and other properties (calibration). The required calculations are complex and finding a global, optimal solution for the equations involved can take a long time. In order to find some good solution in reasonable time, a numeric optimization technique called "bundle adjustment" is commonly used to refine the calibration result iteratively.

Bundler (<http://phototour.cs.washington.edu/bundler>) is an open source SFM software that implements these ideas (Snavely, Seitz and Szeliski 2006). The software can estimate camera parameters and positions, create a sparse scene reconstruction incrementally (i.e. a few images at a time, using bundle adjustment for optimization) and prepare the output for further processing into a dense 3D point cloud with another software. The parameters estimated by Bundler can also be used to remove radial (lens) distortion from the input images.

3.3 Dense 3D Point Cloud Reconstruction

Running SFM software will result in the required camera model(s) and a sparse 3D point cloud that gives a good impression of the 3D geometry (Fig. 4) but is not sufficient for a detailed, realistic reconstruction.

There are a number of different approaches for generating a dense model from the image data and estimated camera parameters (Fig. 4; a comparative study can be found online: <http://vision.middlebury.edu/mview>; see also Seitz et al. 2006). One of the most accurate and complete algorithms is Patch-based Multi-View Stereo (PMVS; Furukawa and Ponce 2007). One of PMVS' advantages is that it preserves only *rigid* structure. E.g. pedestrians walking in front of a monument will not show in the final result. PMVS is also robust against differences in image colours due to exposure settings, white balance or lighting conditions. An open source implementation of PMVS is available (<http://grail.cs.washington.edu/software/pmvs/>).

An alternative approach that is frequently used but will not be pursued further here, is to merge a series of depth maps generated from stereo image pairs (El-Hakim et al. 2008). This method is available via the Epoch/ARC 3D Webservice (<http://homes.esat.kuleuven.be/~visit3d/webservice/v2/>).

3.4. Surface Reconstruction

Dense point clouds can give a very good impression of the 3D scene, but the model will inevitably dissolve into individual points at a certain scale. The reconstructed points therefore need to be connected to form a network of triangles which approximates the shape of the original, continuous surface (Fig. 4). Methods for surface reconstruction can roughly be divided into two classes: sculpting methods that start with a convex hull of the entire point cloud and then "chip away pieces" until the actual surface has been reached, and "region-growing" methods that start with a minimal triangulation and then keep adding new triangles to the model.

Some of the most popular algorithms used in surface reconstruction are Poisson Reconstruction (Kazhdan, Bolitho and Hoppe 2006), Marching Cubes (Lorenson and Cline 1987), Ball Pivoting (Bernardini et al. 1999), Power Crust (Amenta, Choi and Kolluri 2000), Tight Cocone (Dey and

Goswami 2003), or even simple Delaunay triangulation. Unfortunately, most of them have been designed to reconstruct objects that can be scanned or photographed from (almost) all sides and represent "water-tight" bodies, such as sculptures, free standing buildings or architectural fragments.

Archaeological excavation trenches, however, typically show partly uncovered structures, artefacts and deposits with no visible boundaries. Most surface reconstruction algorithms react to these by creating spurious triangles that necessitate some manual cleaning of the result. The only exception is Delaunay Triangulation, which will however fail to reproduce any concave details. Despite this minor problem, we have found Poisson Reconstruction to be a universal and fast reconstruction algorithm that will give good results in almost all cases, after manual deletion of superfluous triangles.

As opposed to laser scan data, 3D point clouds extracted from images are of variable density, which can be a challenge for some algorithms, so some additional pre or post processing of the data may be necessary. The open source software MeshLab has a very complete set of interactive tools that covers every step of the process, from data cleaning and smoothing to surface reconstruction and colour transformation to finally obtain a photo-realistic model.




Fig. 4: Basic steps in image-based 3D modelling. Upper left: SFM derived sparse point cloud. Upper

right: densified point cloud. Lower left: surface reconstruction via triangulation. Lower right: transfer of image colour data.

4. Practical Considerations

Even though the methods discussed here share the property of being robust and requiring little in the way of a data acquisition strategy, they still do have certain requirements that one needs to be aware of.

When considering the use of image-based reconstruction techniques and weighing them against other possibilities, such as classic photogrammetry or laser scanning, specific strengths and limitations of each approach should also be taken into account.

4.1 Photographic Strategy

Purposefully adjusted photographic practice is the most important requisite for successful image-based 3D reconstruction. The better the input data, the better the 3D model will be.

The majority of images should be taken walking around the object, not by pivoting one's own body around a fixed point. After all, the first step in the chain is called "Structure from Motion": If there is no motion (either of the camera around the object or the object past the camera), then very little structural information can be extracted.

It is advisable to take lots of overlapping images, with no more than about 25-30 degrees of angular difference between them. Most importantly, limitations in model detail (data gaps) are generally a function of coverage, not image number. Taking many images from the same location and using the same viewing angles and/or zoom settings will not improve the reconstruction's density; but using more varied angles, view points and zooms will.

Images must never be cropped. This will break the correspondence between focal length and pixel number, thus suggesting a different, wrong scale to the SFM software. But apart from this, current CV algorithms are very good at dealing with differences in image resolution, exposure or lighting conditions. They are also designed to work on raw input images, so extensive pre processing, such as edge enhancement or contrast stretching can actually be counter productive.

4.2 Strengths and Weaknesses

To better assess the role that CV techniques may play in archaeology, it is useful to compare them with the two other, most frequently used approaches: photogrammetric modelling and terrestrial, close-range laser scanning (other scanning techniques, such as structured light, may be based on different physical principles but have very similar properties in practice).

Photogrammetric modelling is an approach that relies on a manual workflow, with a great degree of human judgement and explicit quality control involved at every step. Photogrammetry methods that work in 3D space are parametric, which means that a number of important measures has to be supplied by the operator, such as camera and lens properties or a digital elevation model. They work with relatively few input images and few, explicit correspondences between image features. The methods are task-specific (digital elevation extraction, orthophoto production, projective photogrammetry, etc.)

rather than universal. Complete 3D models must be built by merging smaller components.

By comparison, CV methods rely on massive input datasets consisting of many images and feature correspondences. Quality is assured through high numbers of statistical correlations. They are designed to be non-parametric, requiring only minimal knowledge of the input data, automatic and robust. The approach is, as the name implies, very universal. Once a complete 3D reconstruction has been achieved, task specific products, such as elevation models or planimetric images, can be derived as sub-products (Fig. 5).

While we believe that the CV approach is in many ways superior to Photogrammetry, fits the modern computing paradigm of exploiting massive data sets stochastically, and will therefore eventually supersede these classical methods; the decision between CV and laser scanning is much less obvious. Both approaches have characteristic strengths and weaknesses and cannot be substituted one for the other in all situations.

Typical laser scanners are expensive and slow (compared to cameras) devices that require substantial investment both before and during the field work. However, they do offer considerable return on investment in the form of dense, very regular data that is ideal input for surface reconstruction and highly detailed 3D models. For some applications, the high level of control over the scanning process and the well-known properties of the devices regarding precision and accuracy may also be of considerable value. CV methods, by contrast, produce data of variable density that may be harder to process. However, some algorithms may have additional desirable properties, such as PMVS' ability to reconstruct rigid 3D geometry.

It is also somewhat harder to assess the achievable level of detail for an image-based reconstruction. While the precision is theoretically limited only by the natural boundaries of optical laws, the accuracy is related to the pixel scale. While it can generally be stated, that e.g. an object with a width of 400 cm, photographed with a camera that has 4000 pixels horizontal resolution, will result in an accuracy of 1 mm, a reliable assessment would have to take into account variable scale across the image and overlap between images. There is, however, no grounds for assuming that image-based models are by their very nature less accurate or detailed than laser-scanned ones. After all, both are optical approaches governed by the same physical laws. Indeed, a recent comparison of both techniques, modelling the same object, and published by El-Hakim et al. (2008), has shown the differences to be absolutely negligible.



Fig.5: A side-by-side comparison of the photographic coverage (left) and the extracted model (right; calculated at about half the information density of the photograph) shows the perspective correction of the perfectly orthogonal 3D model.

4.3 Limitations

As has been explained above, the limits of the method are not its precision or accuracy. Modern digital cameras allow for very high resolution images and more detail can always be achieved by covering an object in consecutive stages.

However, processing such large amounts of data can be computationally expensive and while computing time may be less of an issue if calculations are run over night, memory space does represent a hard limit to the level of detail that can be achieved. The number of images is less of a factor, as the algorithms involved only look at small numbers of them at a time, but individual image resolution can drive memory usage up rapidly. We found that with 2 GB of RAM, images of dimensions 2800x2100 could still be processed, claiming 90% of available system resources (the biggest memory bottleneck was the SIFT processing stage). With 8 GB of RAM, 10 megapixel images could be processed without difficulty, memory usage peaking at about 50%.

Demands on both processing time and memory space, however are set to decrease, even as computer hardware becomes ever cheaper and more powerful. The reasons for this are more efficient algorithm designs and an increasing tendency to transfer complex mathematical calculations to the massively parallel processing units of modern graphics cards.

However, the input images themselves may also impose some limitations on what can be achieved. It is well known that scenes with fine, multiple overlap structures (such as the leaves in a tree) will lead to bad reconstruction results, as do shiny and reflective surfaces. In both cases the problem is that the movement patterns will be too complex for the SFM algorithms to cope with. The feature extraction algorithms also have their limitations (Morel and Yu 2009). In particular, images with little discernible structure can be a problem, as well as repetitive patterns, which may lead to arbitrary correspondences.

We have found that SFM algorithms can deal with minimal overlap between images, but PMVS cannot. So while it may be possible to extract just enough points for e.g. interpolating an elevation model, detailed 3D models require sufficient photographic coverage.

4.4 Model Animation and Display

An initial specification for the Weymouth exhibition asked for a 3D model of the mass grave that could be manipulated and explored by visitors (Fig. 6). We were to place a computer monitor and mouse on a table and leave a 3D viewing application running the model for use by passers-by. A Ubuntu Linux desktop running ParaView 3.6.2 (<http://www.paraview.org>) achieved this goal well; the original PLY format file could be loaded and displayed full screen, and in real colours, to be controlled as the user saw fit. Paraview supports a red-cyan anaglyph stereo output, so we were able to supply suitable glasses in order to provide a sensation of depth to the models.

Testing demonstrated, however, that the planned use of ParaView in an interactive environment was not going to be suitable for our exhibition. When used in non-stereo mode, the application was responsive and the model could be manipulated without too much care. The control interface was far from intuitive, however, and although it could be used effectively by users who already had an intimacy with both the software and the archaeology, we had justifiable concerns that exhibition visitors new to both fields, especially children, would find the controls overly confusing. This was most evident when the model was accidentally flipped upside-down and would need to be righted before it made sense again. A second problem came when attempting to produce anaglyph output; despite running the application on a modern computer graphics workstation, we lacked the computing power to smoothly render movement in this mode. When manipulating the model in anaglyph mode, the interactive display would show a greatly reduced resolution level before rendering completely once the model was still; this feature of ParaView is controlled by the LOD (Level Of Detail) Parameters, but we were unable to produce a setting on our machine that provided a model that was viewable during a smooth and easily controlled manipulation.

A second display mode was devised that would see an animation of the model playing on loop throughout the course of the Weymouth exhibition. This would allow us to control the viewing angles in order to highlight the most interesting parts of the model and produce a smooth animation. ParaView was used again to define a series of key frames from which the animation, itself just over a minute long, was constructed. The same animation was rendered twice; once in full colour, and once in anaglyph mode. The available rendering options within ParaView were sufficient for our needs. We chose a rendering resolution of 1920 x 1080, running at 30 frames per second. These videos could then be displayed on monitors throughout the exhibition, as well as being easily provided to the client and disseminated via the Internet.



Fig. 6: A view of the Weymouth mass burial 3D model as presented to the public.

Summary

So far, we have looked at CV mainly as a set of tools for generating rapid and accurate 3D models that can be used for presentational purposes. There are, of course, many more potential uses for this technology in archaeology. Current research at Oxford Archaeology focuses on integrating efficient georeferencing procedures into the workflow, so that the 3D models can be made true to scale and integrated with spatial data from other sources, such as topographic (digital theodolite, GPS) or geophysical surveys. This work is also planned to provide the foundations for the production of further useful outputs, such as orthophotos (Fig. 7), model cross-sections and detailed elevation models. Beyond these, CV algorithms have uses in image database tasks such as classification, object detection, matching and retrieval that still await wider exploitation for archaeological purposes.

As regards the Weymouth case study, the power of current CV algorithms was demonstrated in the form of a detailed model, derived from a small series of images that were not taken with 3D reconstruction in mind. Considering this the results were extremely promising and were successfully used in adding to the excavation project's outreach and communicating its value to the general public, whose appetite for well-presented archaeology, fuelled by TV programmes and other media, only keeps growing. In the future, better integration into ongoing project work and adoption of optimized image taking strategies by field staff will be key priorities, along with rapid digital dissemination of the results.

The different outputs which can be created from the basic 3D model, including video displays and still images for use on web sites and display boards, in lectures and in publications, make image-based reconstruction a versatile, powerful and cost effective tool in disseminating archaeological discoveries to a modern audience. In keeping up with the digital media age, technologies such as 3D PDF bring interactive, 3D site documentation within grasp.

The greatest strength of CV surely lies in its flexibility: the general availability of affordable yet powerful digital cameras and the option to process data on demand, suggest that CV techniques are set to enrich every archaeologist's toolbox. Nothing speaks against taking some extra pictures in the field and thus creating the foundation for a full 3D reconstruction of an object or site at any later point. But existing image archives, ranging from satellite and areal images to site records and historic photographs also provide virtually unlimited resources for further exploration.




Fig.7: A virtual top-down view of the Wemouth mass burial. Photographing such views in the real world requires considerable technical effort.

References

Amenta, Choi and Kolluri 2000: Nina Amenta, Sunghee Choi, Ravi Krishna Kolluri: The Power Crust, Unions of Balls, and the Medial Axis Transform. *Computational Geometry: Theory and Applications*, Vol. 19, 2000, pp. 127-153.

Anderson 2010: Michael Anderson, Precision Recording of Pompeian Standing Remains via Stitched Rectified Photography. *Making History Interactive*, 37th. Proceedings of the CAA Conference (Archaeopress).

Bernardini et al. 1999: Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Claudio Silva, Gabriel Taubin. The Ball-Pivoting Algorithm for Surface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, Vol. 5, 1999, pp. 349-359.

Campana, Remondino 2008: Stefano Campana, Fabio Remondino. Fast and Detailed Digital Documentation of Archaeological Excavations and Heritage Artifacts. *Layers of Perception*. 35th Proceedings of the CAA Conference (Dr. Rudolf Habelt GmbH, Bonn), pp. 36-42.

Dey and Goswami 2003: Tamal K. Dey, Samrat Goswami. Tight Cocone: A water tight surface reconstructor. Proceedings of the 8th ACM Symposium on Solid Modeling Applications, 2003, pp. 127-134.

El-Hakim et al. 2008: Sabry F. El-Hakim, Fabio Remondino, Lorenzo Gonzo, Francesca Voltolini. Effective High Resolution 3D Geometric Reconstruction of Heritage and Archaeological Sites from Images. *Layers of Perception*. 35th Proceedings of the CAA Conference (Dr. Rudolf Habelt GmbH, Bonn), pp. 43-50.

Furukawa and Ponce 2007: Yasutaka Furukawa, Jean Ponce. Accurate, Dense and Robust Multi-View Stereopsis. 2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2007).

Hartley and Zisserman 2004: Richard Hartley, Andrew Zisserman. *Multiple View Geometry in Computer Vision*, 2nd ed. (Cambridge University Press).

Kazhdan, Bolitho and Hoppe 2006: Michael Kazhdan, Matthew Bolitho and Hugues Hoppe. Poisson Surface Reconstruction. *Eurographics Symposium on Geometry Processing 2006*.

Lorensen and Cline 1987: William E. Lorensen, Harvey E. Cline. Marching Cubes: A high resolution 3d surface construction algorithm. *Computer Graphics*, Vol. 21, Nr. 41, 1987.

Lowe 2004: David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60, 2 (2004), pp. 91-110.

Mikhail, Bethel and McGlone 2001: Edward M. Mikhail, James S. Bethel, J. Chris McGlone. *Introduction to Modern Photogrammetry* (Wiley).

Morel and Yu 2009: Jean-Michel Morel, Guoshen Yu. A New Framework for Fully Affine Invariant Image Comparison. *SIAM Journal on Imaging Science*, Vol. 2, Nr. 2, 2009, pp. 438-469.

Seitz et al. 2006: Steven M. Seitz, Brian Curless, James Diebel, Daniel Scharstein, Richard Szeliski. A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms. 2006 IEEE Computer

Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), Vol. 1, pp. 519-526.

Snavely, Seitz and Szeliski 2006: Noah Snavely, Steven M. Seitz, Richard Szeliski. Photo Tourism: Exploring image collections in 3D. ACM Transactions on Graphics (Proceedings of SIGGRAPH 2006).