

Seeing into the Past: Creating a 3D Modeling Pipeline for Archaeological Visualization

P. Allen S. Feiner A. Troccoli H. Benko E. Ishak B. Smith
Department of Computer Science, Columbia University, New York, NY

Abstract

Archaeology is a destructive process in which accurate and detailed recording of a site is imperative. As a site is exposed, documentation is required in order to recreate and understand the site in context. We have developed a 3D modeling pipeline that can assist archaeologists in the documentation effort by building rich, geometrically and photometrically accurate 3D models of the site. The modeling effort begins with data acquisition (images, range scans, GIS data, and video) and ends with the use of a sophisticated visualization tool that can be used by researchers to explore and understand the site. The pipeline includes new methods for shadow-based registration of 2D images and temporal change detection. Our multimodal augmented reality system allows users wearing head-tracked, see-through, head-worn displays to visualize the site model and associated archaeological artifacts, and to interact with them using speech and gesture.

1. Introduction

Archaeology is a destructive process. Analysis of a site requires the dismantling of structures and the displacement and removal of findings, such as tools, pottery and bones. Therefore, it is important to keep an accurate record of each site as it is being excavated. Our work focuses on developing new tools and methods to model and visualize historic and archaeological sites.

In this paper we address several problems. The first is how to digitally record the progress of an ongoing archaeological excavation using different information sources: 3D laser scans, digital video, panoramic photography and GIS (geographic information systems) data. Our technical goal is to complement existing field methods with these technologies, as shown in Figure 1. Building 3D models from range scans and images is time-consuming and difficult, usually involving much manual effort. In a typical 3D modeling pipeline, geometry is acquired in the form of range scans that need to be registered together in a common coordinate system. Additionally, images provide texture information. To build a texture-mapped model, the images



Figure 1. Members of our team record exact 3D coordinates using a total station surveying instrument (upper left), capture video of an excavation using a scanner and digital camera (lower left), and scan a portion of the excavation site using a time-of-flight laser range scanner (right).

and the geometry have to be registered (i.e., a mapping from model space to image space has to be computed). Part of our research focuses on developing new methods for reducing model-building time and improving accuracy [2]. In particular, texture mapping can pose difficult problems in providing complete and accurate coverage of a complicated model. Our method uses knowledge of shadows and sun position during the image acquisition process to correctly register the imagery with the constructed 3D model. But our interest goes beyond building a 3D model; we also want to record and keep track of the changes to the site as the excavation proceeds. This problem also poses interesting technical challenges. Scanning the entire site every day is not desirable or practical. Instead we would like to acquire enough information to track changes, so as to be able to build a 3D model that represents a snapshot of the site at any given point in time.

A second problem we address in this paper is the development of a visualization tool that will enhance the post-excavation analysis and review process. For this purpose, we have built a multimodal augmented reality system for




Figure 2. Our 3D modeling and visualization pipeline. We start by building a textured 3D model using range scans and images, which we enhance with contextual information in the form of panoramic images, video, and GIS data. This context-rich model is then used as input to our multimodal augmented reality application.

viewing the 3D model and the accompanying contextual information. In our prototype, users wearing head-tracked, see-through, head-worn displays can explore a full-size representation of a virtual site, and inquire about findings and objects using both gesture and speech. Additional information in the form of still images, video and textual description is also available. Figure 2 describes our modeling and visualization pipeline. Each of the steps in the pipeline is described in the sections that follow.

2. Previous Work

3D modeling. Over the past few years, a number of research teams have been addressing the use of range scans and images to develop 3D models for the virtual preservation of historic sites. Some notable projects include the modeling of Michelangelo’s David and other statues by Levoy et al. [10], the IBM Pieta project of Bernardini et al. [4], the Great Buddha project of Ikeuchi et al. [8], and the virtualization of a Byzantine Crypt by Beraldin et al. [3]. Our goals overlap with the work of these researchers, but also differ in several ways. First, we are interested in recording an archaeological excavation in progress. Second, we want to keep track of changes as the excavation proceeds. And finally, we require that our models serve as a complement to the archaeologists’ documentation, which means sharing the same reference coordinate system. The work of Acevedo et al. [1] is similar to ours with respect to these latter goals. However, they use photogrammetry,

instead of range data from laser scans, to create their 3D model; hence, the technical challenges we each face are quite different. Image-based reconstruction for archaeology has also been addressed by Pollefeys et al. [12].

Visualization. Archaeologists currently use many kinds of written documentation, sketches, diagrams, and photographs to document the physical state of an excavation site, following standard guidelines that have been developed for this purpose. To visualize the data, most archaeologists currently rely on GIS software, such as ESRI’s ArcGIS. Recently, INTRASIS has extended some capabilities of standard GIS systems and functions as a plug-in for ArcGIS. Standard computer-aided design (CAD) systems, such as AutoCAD, are often used for modeling and reconstruction, which is both costly and time-consuming. While both GIS and CAD contain 3D visualization capabilities, most of those systems tend to present layered 2D maps or coarse topographical terrain maps with embedded objects, sketches, and pictures. They do not support additional multimedia, such as audio, video, 3D high-resolution terrain scans and panoramas.

Several research groups have explored immersive 3D visualization for archaeology. For example, the ARHAVE project [1] was developed for use in a CAVE-like environment with projected walls and floor [5]. It consists of a human-modeled environment embedded with virtual icons representing various types of finds and has been used to determine patterns and trends of the objects found on site. In contrast, our system uses a highly detailed computer-

generated model, containing high-resolution textures, making possible a more accurate representation of the actual site. Gaitatzes et al. [7] present virtual reality environments for visualizing temples and public buildings of ancient Greece. Relatively stationary users visualize an archaeological site on a tilted rear-projected desk or in a CAVE-like room. In contrast, users of our system can physically walk around the life-sized excavation site to explore the virtual world from different perspectives.

In the following sections, we present the 3D modeling pipeline and visualization facilities that we have developed. We illustrate them with examples based on data that we acquired on-site at the excavation of a 6th-4th century BC archaeological excavation at Monte Polizzo in western Sicily.

3. The 3D Modeling Pipeline

Geometry, in the form of point clouds, is acquired by a laser range finder, and texture, in the form of photographs, is obtained with a digital camera. Details for each stage are given below. A video of our model can be downloaded from www.cs.columbia.edu/~allen/sicily.avi

Scan acquisition. To model the acropolis at Monte Polizzo we used a time-of-flight laser scanner (Cyrax 2500) to measure the distance to points on the site. Data from the scanner comprises point clouds, with each point consisting of three coordinates (x,y,z) and a value representing the amplitude of the laser light reflected back to the scanner.

Scan registration. Multiple scans are required to completely acquire a site such as the Monte Polizzo acropolis. The resulting point clouds need to be registered together. Typically, the coordinate system of one of the point clouds is chosen as the coordinate system for the model. In archaeology, however, a global site coordinate system is set up from GPS data. A set of control points are accurately measured using (preferably differential) GPS, and are then used to initialize a total station (e.g., a Leica TCR 705 total station, as shown in the top left portion of Figure 1) in the world coordinate system. A total station is a theodolite with an electronic distance meter that is used to measure points of interest, such as the location of findings, rocks, or the terrain contour. To register each point cloud with respect to the site's coordinate system, we use a set of targets that the scanner can automatically recognize, shown in Figure 3. Before taking a scan, we place the targets on the area we plan to cover, and use the total station to measure their positions in the site's coordinate system. Afterwards, we scan the scene at a low resolution to identify and acquire the targets' positions in the scanner's coordinate system, and so solve the 3D-to-3D registration problem. The targets are then removed and a full-resolution scan is acquired.

This technique allowed us to accurately register each individual point cloud with the site's coordinate system.



Figure 3. Modeling the acropolis at Monte Polizzo. Top: Targets have been placed in the area to be scanned. Bottom: Final textured model with panoramic image as background and GIS data.

For our purposes, it proved advantageous over pairwise or multi-scan registration using the iterative closest point algorithm (ICP) for several reasons. First of all, it required no scan overlap. This allowed us to take fewer scans, with greater freedom to choose the scanner position by eliminating the traditional overlap requirement of ICP. We did, however, acquire overlapping scans to minimize unseen regions (holes). In addition, as soon as we finished a scanning day, our model was completely registered. And finally, it allowed us to record changes easily by scanning only the affected areas. We would not have been able to track changes robustly if we had relied on ICP for scan registration, because ICP relies on point-to-point correspondences over an overlapping set of point clouds. If the point clouds represent different states of the site, these correspondences may not exist.

Surface generation. From sets of registered point clouds that represent the state of the site at the same point in time, we generated a triangular-mesh surface, using the VripPack software developed by Curless and Levoy [6]. VripPack outputs the best mesh that fits the point cloud data,

smoothed to account for registration errors.

3.1. Texture mapping

Texture acquisition. In addition to the scanner, we used a Nikon D100 digital camera, mounted on the scanner's case, to acquire texture information. For each scan we acquired, we took a photograph.

Local texture registration. Prior to our trip, we performed a simple calibration to estimate the camera's external and internal parameters. We determined the camera calibration by scanning a flat wall with the room lights off and the camera's shutter open for the eight-second duration of the scan. This provided us with an image of the grid pattern described by the laser as it sampled the wall and the 3D coordinates of each sample. We scanned again from a different distance and angle to acquire more samples. We then segmented the images to obtain the centroid of each grid point, and solved the calibration problem using the 2D-to-3D correspondences just obtained.

Global texture registration. While the local texture calibration procedure provided us with a good estimate of the camera's parameters, we found that our images were slightly misaligned with respect to the complete model. One reason for this is that our calibration was local to the scanner's coordinate system. To texture-map the final model, this local registration had to be transformed to the site's coordinates. Hence, any errors in scan-to-scan registration will also affect the texture registration. In addition, our initial calibration was accurate at the depths at which calibration points had been measured, but not as accurate at other ranges. To solve these misalignments, we developed a new method based on the shadows cast by the sun. Our method performs a global texture registration; it registers the texture with respect to the model's coordinate system, as opposed to the scanner's coordinate system. Since we have the latitude and longitude of the site and the time at which each photograph was taken, we can compute the location of the sun and find portions of the 3D model that should be in shadow. By matching these with the shadows in the image we solve the 2D to 3D registration problem.

Assuming the internal parameters of a camera are known, we find the camera position c with respect to the 3D model: $c = (\phi_x, \phi_y, \phi_z, t_x, t_y, t_z)$. This is a six-parameter rigid body transform that maps a point X_w in world coordinates into its corresponding point X_c in the camera reference frame. The first three parameters (Euler angles) represent the angles of rotation about each of the coordinate axes and form a rotation matrix, $\mathbf{R}(\phi_x, \phi_y, \phi_z) = \mathbf{R}_x(\phi_x)\mathbf{R}_y(\phi_y)\mathbf{R}_z(\phi_z)$. The remaining three parameters are the components of a translation vector \mathbf{t} . Together, they satisfy the following relationship:

$$X_c = \mathbf{R}(\phi_x, \phi_y, \phi_z)X_w + \mathbf{t}.$$

If we knew the correct set of external camera parameters $(\phi_x, \phi_y, \phi_z, t_x, t_y, t_z)$, then an orthographic view of a textured version of the model with the eye looking in the direction of the sun's rays should show no texture representing shadows. However, if the texture is misaligned, such a rendering will exhibit a number of shadow pixels. Our method exploits this idea by searching the parameter space for a point that minimizes the number of pixels representing shadows in the rendered image of the model.

The problem is properly stated as follows. If we let I denote the image to be registered and M the model, then f , our cost function, is defined as

$$f(I_r) = \sum_{x,y \in I_r} shadow(I_r, x, y),$$

where I_r stands for a rendered image of M as seen from the direction of the sun and textured with I using a texture camera with external parameters set to \mathbf{c} , and $shadow(I_r, x, y)$ is 1 if pixel (x, y) of I_r is in shadow, otherwise 0.

Given the initial estimate of the camera position \mathbf{c}_0 that we obtained off-line with our pre-calibration, the problem is to find a point \mathbf{c}_f that minimizes f .

The complete shadow registration process consists of two stages: a preprocessing stage and a minimization stage. In the preprocessing stage, the shadows in the image are found using thresholding, and masked out with a given color. In the minimization stage, simulated annealing is used to search for a global minimum of f , starting from the initial estimate. We applied the algorithm to ten of the texture images we used to create our model. Figure 9 shows screenshots of the final textured model. A detailed analysis and quantitative results can be found in [14].

Texture-map generation. To create the final texture-mapped model, we assigned each mesh triangle one of the available textures. For each vertex v , we first find its valid image set I_v . An image i_k belongs to the valid image set of a vertex if the following three conditions are met: first, the projection of v must be within the boundaries of i_k ; second, the normal n_v of v must be pointing towards the camera of i_k (i.e., the angle between the optical axis of the camera and n_v must be less than $\pi/2$); and, finally, there must be no other point in the mesh M in the line connecting the projection of v in i_k and v (these conditions are mentioned in [13]). We perform the last test using a ray-tracing operation accelerated with an octree. We then compute for every triangle t its valid image set I_t . An image i_k is in I_t if it belongs to the valid image set of each of the triangle's vertices. Finally, from the valid image set of each triangle, we choose the image i_k with the best pixel/area resolution. The final textured model is rendered using hardware-supported projective texture mapping.




Figure 5. Tracking changes. Left: Initial model. The stones to be removed are numbered. Middle: Model after stones have been removed. Right: Detected changes, in light gray and numbered. The careful observer will notice some small gray patches that do not represent removed stones. These are due to small misalignments and smoothing errors.




Figure 4. Incremental changes in the model are detected by casting a ray from the scanner position to the model through point p_i , and finding its intersection point m_i .

3.2. Model change detection

We have also developed a method to track changes to the excavation site. Because of the large size of the Monte Polizzo site, we limited ourselves to a single structure, approximately 10 by 10 meters in area. We first acquired the scans necessary to build a full initial model of the structure. The archaeologists then proceeded to remove some stones and we acquired new scans. We did not scan the entire structure again; instead we scanned only those areas where stones had been removed. Using our change tracking method, we can recreate the site with and without the removed stones, simulating the archaeological process itself. Figure 5 shows renderings of the site model at different stages, where the numbers labels the stones that were removed. The first image shows the site in its initial state, the middle one shows the site after the stones were removed, and the right image shows the detected changes in gray.

During post-processing, we created an initial model M_0

of the structure by registering and merging the initial point clouds together. We then incrementally incorporated the changes to the model. From the initial model M_0 and a point cloud P that was acquired after the stones were removed, we detect the geometry corresponding to the removed stones, delete it, and add the geometry corresponding to the newly exposed surfaces to obtain a new instance M_1 of the model, in the following manner:

1. Let X_p be the position of the scanner when P was taken.
2. For every point p_i in P do
 - (a) Trace a ray R from X_p through p_i (see Figure 4).
 - (b) Find the triangle t_i of M_0 that R intersects.
 - (c) Find m_i , the intersection point of R and t_i .
 - (d) Compute the distance d_i between p_i and m_i .
 - (e) If d_i is greater than a given threshold, and R passes first through m_i and then through p_i , remove all edges of triangle t_i from M_0 .
3. Remove all non-connected vertices from M_0 . Call this the intermediate model M_{im} .
4. Create a mesh M_p from P and use a package such as VripPack to merge M_{im} with M_p to obtain M_1 .

These steps are repeated for all point clouds acquired after the stones were removed. The result is a new model that represents the updated state of the site.

3.3. Adding context information

3D modeling from range scans is a powerful tool, but a 3D model by itself lacks important context information. We can obtain this context by combining data from different sensors. In our pipeline, we combine our 3D model with surveying data, panoramic images, and digital video.

In current archaeological recording methods, a total station is used to lay out a site grid and record the 3D position of finds. The logged data is then used to make a site plan, typically using GIS software. We can incorporate this data in our modeling pipeline to add important and meaningful information. For example, by displaying the site grid, archaeologists that were not at the site can relate written reports to the 3D model. The GIS data and the 3D model are, by design, in the same coordinate system and no special registration is required.

Additional context information is provided by panoramic images. In the field, we acquired a complete panorama from a set of thirty eight images using a Kaiden QuickPan III spherical panoramic head. The final cylindrical panorama was created with PhotoStitch, a photo-stitching package. For registration purposes, we recorded the position of the camera using the total station. (We did not record the camera's orientation, which we had to find manually.) Since the camera was leveled, finding the correct rotation is a one-degree-of-freedom search (rotation about the camera's y axis), which we performed manually by rotating a textured-mapped cylinder around the model.

Finally, we used a color digital video camera to capture moving imagery of the excavation in progress, recording the position of the video camera with the total station.

4. Visualization for Archaeology

Archaeological excavations are projects that involve many experts and usually last for years. Most of that time is actually spent away from the excavation site in preparation for each excavation season, or in analysis of data that has been excavated during the previous seasons. In addition, as the site is excavated, the layers are removed without ever being able to physically recreate them, and most of the excavated material is stored in a museum repository that is not easily accessible for analysis. Our system is intended to explore how these problems might be solved. By digitally recording the excavation, we allow archaeologists to remotely visualize the site at will. We combine a 3D layered terrain model of the site, with models of excavated objects that are correctly located in the site model.

In our visualization environment, shown in Figure 8, multiple users, each wearing a head-tracked, see-through, head-worn display (Figure 6), can simultaneously examine the site model. We believe that this approach has several benefits compared to CAVEs [5] and similar environments, in which one or more users occupy a room whose walls (and, possibly floor and ceiling) are projected with imagery computed for the perspective of a single tracked user:

- Multiple users can each have a correct 3D stereo view of the same or a different part of the the environment, depending on the user's location.




Figure 6. A user, wearing a head-tracked, see-through head-worn display, an instrumented glove, and a hand tracker.




Figure 7. Live video shot showing the visual representation of four SenseShapes that are used in our system; two are attached to the user's head (representing the user's view frustum) and the other two to the hand (a pointing cone and a grabbing sphere).



Figure 8. Our augmented reality visualization system. A user in our lab explores a full-size model of the site and selects one of the findings to examine in detail. This image is taken through a live, tracked, video see-through display.

- The active area can be large enough (limited only by the tracking technology), that the excavation site can be explored at full size, as users walk around it.
- See-through displays allow users to see each other and interact without the need of virtual avatars.
- While head-worn displays of the sort we currently use have a lower pixel count (800×600) than one wall of a typical current projection system (1024×768 or 1280×1024), the otherwise disadvantageous smaller field of view of the head-worn display means that each pixel subtends a smaller visual angle.

Our visualization environment is intended for many applications, including post-excavation analysis, educating archaeology students about excavation procedures, and training new personnel prior to the excavation season.

Each user is able to view a 3D representation of the terrain. In addition, small archaeological finds are placed around the site at the exact locations of their discovery, each labeled with its name and description, as shown in Figure 8. A 360-degree panorama surrounds the model to provide an immersive virtual experience. As a user walks around the virtual site, she can inquire about situated objects through a combination of speech and gesture, or navigate using a wireless mouse. Additional information about specific objects is available in the form of still images, movies, and textual descriptions that can be presented on the head-worn display.

Digital images, taken on site, make it possible for archaeologists to visualize the site's overall structure, as well as the geometry and texture of individual stones, both those removed during the excavation and those still in place. Movies allow archaeologists to witness the excavation process. They also provide an invaluable walkthrough perspective of the site, making it possible for archaeologists to experience the excavation both visually and audibly through the eyes of the camera-person. Archaeologists also have access to a textual database of information about the various trenches, objects, and excavations on site. This information is available both in a tabular fashion in a fixed portion of the display, as well as in situ, labeling the objects found in the 3D environment.

Our environment uses a multimodal interaction engine that combines speech input with glove-based gestures to facilitate easy and natural object selection. We are using IBM ViaVoice V.10 for speech recognition and EssentialReality P5 gloves with our own gesture recognizer. Currently we can differentiate between three gestures (point, grab, thumbs up) and an idle state. To help disambiguate references to closely located objects, we have used SenseShapes as a primary mechanism for interacting and selecting the objects in the environment [11, 9]. *SenseShapes* (Figure 7) are

volumetric regions of interest that are attached to parts of a user's body (e.g., the head and a hand) to provide statistical information about the user's interaction with objects. The statistics describe, for any selected time interval, how long an object has intersected a SenseShape's volume, the number of times the object has entered and exited the volume, the distance of the object from the volume's origin, how much of the object is visible from a key point within the volume, and how close the object is to the volume's center line. The system keeps a complete history of object intersections with the SenseShapes in an in-memory database, thus making it possible to aggregate the statistics over time. This effectively allows the system to formulate higher level queries that express requests such as "Return all objects of type U and statistics associated with them that were in the hand-picking cone and in the user's view cones that happen to have intersected the hand-picking cone between time V and time W ." The returned object list is then sorted based on a weighted average of the object statistics. Furthermore, by combining this hypothesis list with the output of the speech recognizer and the gesture recognizer, the system can make fairly sophisticated multimodal decisions.

5. Summary and Conclusions

We have described an integrated 3D modeling and visualization pipeline for archaeology, which we have applied to digitally recording an excavation in progress at Monte Polizzo, Sicily. The area of 3D modeling for cultural heritage preservation is evolving rapidly and we believe that new tools such as these will be an important resource for future archaeological research. Our work is unique in a number of ways, because it 1) integrates a variety of different data sources: range scans, images, GIS data, and video, 2) develops a new technique for image-to-model-base registration based on the shadows cast by the sun, 3) includes a temporal change tracking component, and 4) presents a novel set of multimodal augmented reality visualization tools for archaeology that can be used with this and other models.

Acknowledgments

Lynn Meskell and James Conlon assisted with the field work at Monte Polizzo. Ian Morris and the Stanford Archeology Center graciously provided access to and images and data from Monte Polizzo. Sajid Sadi and Avanindra Utukuri provided support with the Essential Reality P5 glove. This research was funded in part by ONR contracts N00014-99-1-0394, N00014-99-1-0683, N00014-04-1-0005, and N00014-99-1-0249, NSF grants IIS-0121239 and IIS-00-82961, and gifts from Mitsubishi Electric Research Labs, Microsoft Research, and Alias Systems.



Figure 9. Screenshots of our 3D model with the panorama in the background.

References

- [1] D. Acevedo, E. Vote, D. H. Laidlaw, and M. S. Joukowsky. Archaeological data visualization in VR: Analysis of lamp finds at the Great Temple of Petra, a case study. In *Proc. IEEE Visualization '01*. IEEE Computer Society, 2001.
- [2] P. K. Allen, A. Troccoli, B. Smith, S. Murray, I. Stamos, and M. Leordeanu. New methods for digital modeling of historic sites. *IEEE Comp. Graph. and Applic.*, 23(6):32–41, 2003.
- [3] J.-A. Berladin, M. Picard, S. El-Hakim, G. Godin, G. Valzano, A. Bandiera, and D. Latouche. Virtualizing a Byzantine crypt by combining high-resolution textures with laser scanner 3D data. In *Proc. VMMS 2002*, pages 3–14, September 2002.
- [4] F. Bernardini, H. Rushmeier, I. M. Martin, J. Mittleman, and G. Taubin. Building a digital model of Michelangelo's Florentine Pietà. *IEEE Comp. Graph. and Applic.*, 22(1):59–67, Jan/Feb 2002.
- [5] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proc. SIGGRAPH 93*, pages 135–142. ACM Press, 1993.
- [6] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proc. SIGGRAPH 96*, pages 303–312. ACM Press, 1996.
- [7] A. Gaitatzes, C. D., and R. M. Reviving the past: Cultural heritage meets virtual reality. In *Proc. Conf. on Virtual Reality, Archeology and Cultural Heritage*, pages 103–110, 2001.
- [8] K. Ikeuchi, A. Nakazawa, K. Nishino, and T. Oishi. Creating virtual buddha statues through observation. In *IEEE Workshop on Applics. of Comp. Vision in Architecture*, 2003.
- [9] E. Kaiser, A. Olwal, D. McGee, H. Benko, A. Corradini, X. Li, P. Cohen, and S. Feiner. Mutual disambiguation of 3D multimodal interaction in augmented and virtual reality. In *Int. Conf. on Multimodal Interfaces*, pages 12–19. ACM Press, 2003.
- [10] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk. The digital Michelangelo project. In *Proc. Siggraph 2000*, pages 131–144, 2000.
- [11] A. Olwal, H. Benko, and S. Feiner. SenseShapes: Using statistical geometry for object selection in a multimodal augmented reality system. In *IEEE and ACM Int. Symp. on Mixed and Augmented Reality*, pages 300–301. IEEE Comp. Soc., 2003.
- [12] M. Pollefeys, L. Van Gool, M. Vergauwen, K. Cornelis, F. Verbiest, and J. Tops. Image-based 3D acquisition of archaeological heritage and applications. In *Conf. on Virtual reality, Archeology, & Cultural Heritage*, pages 255–262. ACM, 2001.
- [13] C. Rocchini, P. Cignomi, C. Montani, and R. Scopigno. Multiple texture stitching and blending on 3D objects. In *Rendering Techniques*, Eurographics, pages 119–130. Springer-Verlag Wien New York, 1999.
- [14] A. Troccoli and P. K. Allen. A shadow based method for image to model registration. In *IEEE Workshop on Image and Video Registration, CVPR*, 2004.