Reconstructing the Thera Frescoes: Automated Digitization and Matching of Fragments

Executive Summary

The archaeological site of Akrotiri on the island of Thera (modern-day Santorini) has proven a treasure trove of information about Aegean civilization and culture. Among its most valued artifacts are wall paintings (frescoes), which have been preserved in the volcanic ash since the sixteenth century BCE. The frescoes, however, are typically recovered in fragments of a few centimeters to a few tens of centimeters in length, and reconstructing complete wall sections from the fragments occupies a major portion of the effort at Akrotiri.

We propose to build a system that will assist archaeologists and conservators by digitizing excavated fragments and automatically proposing matches on the basis of color, 3-D shape, and other cues. Such systems have been researched before, for objects including frescoes, pottery shards, and inscribed marble slabs. For this project, however, our overriding goal is to leave an ongoing legacy: a combination of equipment and computer software that is sufficiently practical and easy to use that it will become an everyday part of excavation activities at Akrotiri. The system will have three major components:

- **Acquisition equipment.** Beginning with commercially-available color and 3-D scanners, we will build a hardware configuration for capturing the color texture of each fragment, together with its thickness and 3-D shapes of the edges. We will write custom software to control the scanners, ensuring that capture is easy (may be performed by non-computer-experts) and fast (no more than a few minutes per fragment). All captured data will be stored together with semantic information about the fragment, as well as any annotations provided by archaeologists and conservators.
- **Matching algorithms.** Building on recent research results, we will develop a system for computing candidate matches between fragments, and proposing them to the user for verification. We will explore different cues for performing this matching, including the 3-D shape of edges, fragment thickness, color and edges on the surface, any user-supplied annotations, and arbitrary combinations of different cues.
- **User interface.** We will build a computer program that allows users to move fragments or groups of fragments on the screen, easily evaluating the proposed matches. The software will keep track of all information associated with each fragment, as well as with each match. The software will also include features such as revision control and support for concurrency, allowing multiple users to work either simultaneously or serially.

On the basis of our past experience, we believe that such a system will successfully allow for a more rapid and more complete reconstruction of the wall paintings, with less handling of the physical fragments. The project will be carried out in cooperation with archaeologists and conservators at Akrotiri, led by the director of excavations, Christos Doumas.

We propose to carry out the work beginning early in 2007, with major development complete by the beginning of the summer. Software development will be carried out by researchers and graduate students at Princeton University, and will be coordinated with a course taught at Princeton during the spring semester. This proposal describes in detail the components of the hardware and software systems that we will develop, together with our educational plan.

1 Introduction

The island of Thera (modern-day Santorini) was the site of one of the most advanced civilizations of the Aegean world, displaying extraordinary engineering and artistic prowess. A massive volcanic eruption (dated by recent estimates [3] to between 1627 and 1600 BCE) buried the island in ash, preserving an entire city complete with pottery, buildings, and, most significantly, wall paintings (frescoes). These were shattered by the force of the eruption, the accompanying earthquakes, and the passage of time, yet provide unique insights about Aegean civilization and culture (Figure 1). Since the beginning of excavations in 1967 at the site called Akrotiri, the frescoes have become the most famous products of the Thera excavations [2]; still, reconstructing complete wall panels from the thousands of fragments, ranging in size from a few centimeters to a few tens of centimeters, occupies a majority of the archaeological effort expended at Akrotiri.

There is substantial room for technology to assist in matching fragments—solving the "jigsaw puzzle"—as well as in the further study of the assembled frescoes. Early research in this direction has not only developed algorithms that can help in the assembly of fragments, but also demonstrated that at least some of them were painted using a small number of stencils [9, 8]. Furthermore, these stencils included hyperbolas and linear spirals—geometric shapes that neither the Minoans nor their contemporaries were previously thought to be able to construct.

Automated assembly work has been investigated in the context of other archaeological projects as well. Much of this work has been done on *potsherds*—fragments of pots [4, 1, 12]. While pottery assembly is similar in spirit to fresco assembly, only some of these matching techniques carry over. The pottery considered is generally thrown on a wheel, and the axial symmetry around the center of the pot is used to guide the matching.

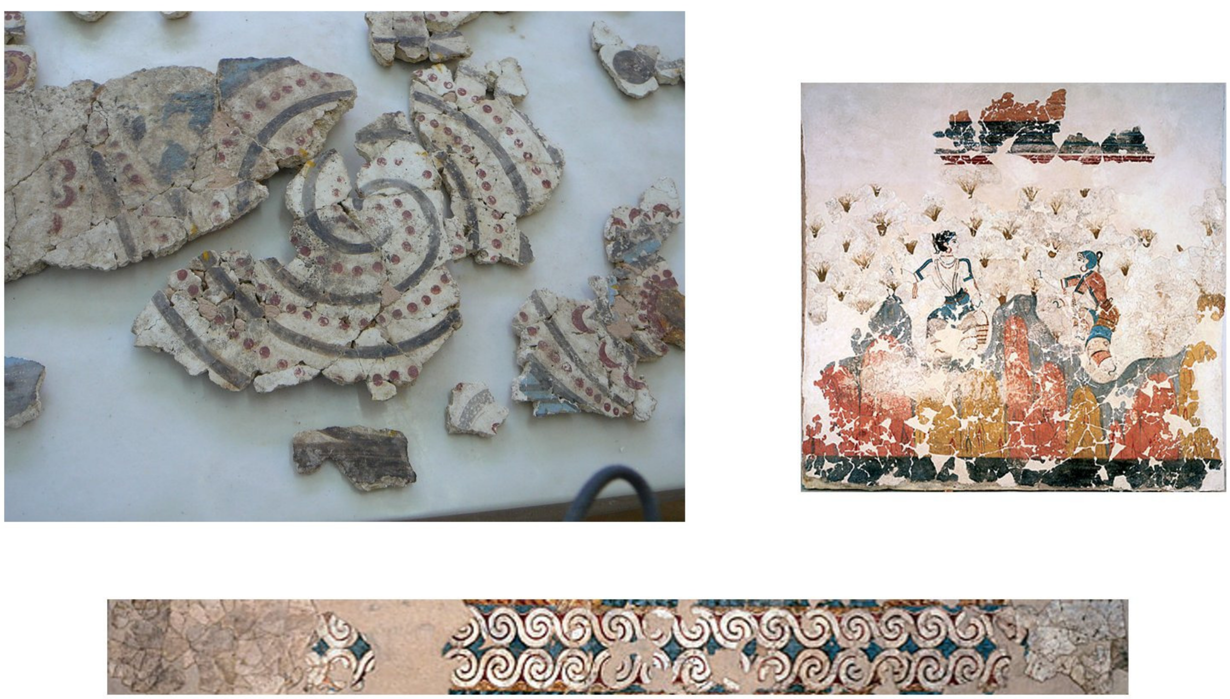

Figure 1: *Frescoes from Akrotiri. Clockwise from top: A spiral in course of assembly; Saffron Gatherers (Xeste 3, Room 3a); Frieze with Spirals (Xeste 3, Room 2).*

Another related project is the Forma Urbis Romae project [6], conducted at Stanford University, to acquire 3-D scans of this second-century marble map of Rome. Research into the assembly of these fragments is still ongoing, though already the research effort has yielded 20 new matches that had not been found by decades of manual labor [5].

Computer technology has also been used to help archaeologists visualize and documents sites in new ways. One of the most ambitious of such projects is the ARCHAVE project at Brown University [10], which records the information about finds at the Great Temple in Petra along with their precise locations. The results of database queries can then be visualized in a virtual model of the temple in a completely immersive 3-D environment. This provides many interesting possibilities, such as quickly seeing how lamps of a particular style were distributed around the temple.

Our project for fresco assembly, while modest relative to some previous projects, adopts the ambitious yet realizable goal of producing a system that can be operated solely by non-experts. We will build acquisition and reconstruction systems suited to the specific needs of fresco assembly, incorporating state-of-the-art scanning hardware, color/3-D matching methods, and visualization tools. We hope that these systems become fixtures on this and other archaeological sites, reducing the burden of matching the majority of fragments. In addition, the matching techniques developed for this project will push the state of the art in the field of 3-D geometric processing, while the datasets we acquire may enable longer-term research into analysis of, for example, painting techniques or pigment usage.

We believe that the proposed project is very likely to be successful. Our team includes computer scientists at Princeton University and the National Technical University of Athens, with extensive experience in digital imaging and 3-D scanning for cultural heritage applications as well as in methods for matching on the basis of color and 3-D shape. We have the support of the Program in Hellenic Studies at Princeton University as well as the Thera Foundation, with initial funding commitments from each. The remainder of this proposal describes in detail the components of the hardware and software systems that we will develop, together with our educational plan.

2 Acquisition

The goal of the acquisition system is to capture as much information about each fragment as efficiently and as simply as possible. Certainly, the system must capture a high-resolution color image of the top of each fragment, in order to preserve the painting. However, archaeologists and conservators also use the 3-D shape to find matches, not just the color: the exact shapes of the edges of matching fragments frequently "lock" into place, either along the entire matching edge or, due to erosion, along a part of it. We will therefore capture the 3-D shape of the sides of each fragment, allowing the computer to perform similar kinds of matching and to verify matches without recourse to the original artifacts. In addition, we will capture the 3-D shape of the *bottom* of each fragment. This information provides a number of less obvious benefits: it gives information about the thickness of each fragment, which can be used to aid matching; together with the edge and color information it provides an essentially complete 3-D representation of each fragment for archival and study purposes (we assume that the painted surface is nearly flat); finally, this data can simplify the technical task of aligning the 3-D and 2-D information.

2.1 Acquisition of 3-D Shape

We will use a commercial 3-D light stripe scanner to capture the complete shape of each fragment, including the thickness and the detailed profiles of the edges. These scanners work by shining a stripe of light, produced by a very low-power laser, onto the surface of the object. A camera integrated with the scanner, but oriented at an angle with respect to the light source, captures the image of the stripe (Figure 2). The

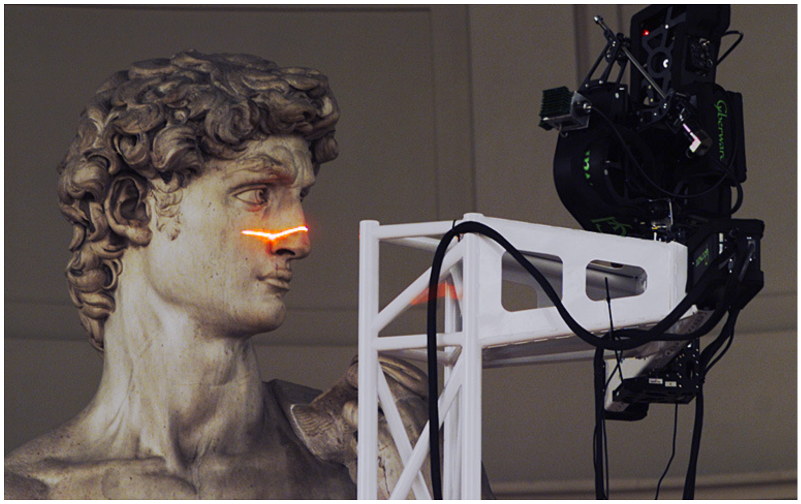

Figure 2: *Light stripe triangulation scanner being used to scan the David, during the course of the Digital Michelangelo Project [7]. A stripe of light is projected onto the surface, and a camera observes the stripe at an angle. The apparent deformation of the stripe is due to differences in depth, and the principle of triangulation is used to obtain 3-D point positions.*

principle of *triangulation* is then used to determine the 3-D position of each point along the stripe. The mathematical basis of triangulation is that the position of the light stripe (i.e., a plane of light) is known, as are viewing rays from the camera to the observed positions of the light stripe. 3-D coordinates may therefore be computed as the intersections of the viewing rays with the light plane. Because only a single contour of 3-D data is captured at once, the stripe is swept slowly along the object to capture one entire side of the object; the object is then automatically rotated on a computer-controlled turntable and the process repeated, allowing the entire surface to be captured.

While there are a few other technologies available for acquisition of 3-D shape, we believe that light-stripe scanners have several benefits:

- The scanner is a *non-contact* system, meaning that the fragment is never touched. This ensures safety and prevents erosion.
- The light source used is very *low-power*—typically a few milliwatts—meaning that the resultant light intensity at surface is comparable to that produced by a standard light bulb at one meter. Moreover, the energy produced by a laser is a single color, meaning that it introduces no harmful ultraviolet or infrared.
- Light stripe scanners return very *accurate* data: typical specifications for 3-D accuracy range from 0.01 to 1 millimeter, depending on the scanner. For this project, we will evaluate scanners with accuracy in the range of 0.1 to 0.4 mm.
- The use of light stripe scanners has been *demonstrated* in archaeological and cultural heritage applications (such as the project to digitize the 1186 existing fragments of the Forma Urbis Romae, a map of 3rd-century Rome carved into marble [6]).

2.2 Acquisition of Color

We are investigating two possible systems for color acquisition, each with possible benefits and drawbacks. The final decision regarding which system is deployed will be made in consultation with all collaborators on the project, especially the archaeologists and conservators, to ensure that it is safe and reliable.

Flatbed Scanner: Our currently-preferred capture method would be to use a digital flatbed color scanner, as commonly used for scanning documents. Such scanners are inexpensive, high-quality, and very high-resolution (tens of pixels per millimeter). Because the fragments would be placed flat-side-down on the scanning bed, it is very simple to operate and to place fragments (though we must ensure that doing so is safe). Moreover, as discussed below, there are several possible configurations for conveniently integrating the color scanner with a 3-D scanner, assuring that the two types of data are registered to each other.

Flatbed scanners do have certain drawbacks. First, there is a question of light exposure from the scanning process. However, given the low light intensities used by modern scanners and the brief length of the scanning process, the total light exposure due to scanning is comparable to significantly less than an hour of indoor illumination, or approximately 1 second of direct sunlight [11].

A second potential concern is whether irregularities in the fragments will cause portions of the surface to be out of focus. However, in our initial investigations, we have found that scanners maintain adequate focus for distances approaching a centimeter from the scanning bed. Therefore, we believe that in all cases we will be able to scan the top surface of each fragment without problems.

A more significant challenge involves keeping the glass surface of the scanner clean and unscratched, given the dusty environment and repeated placing and removal of the fragments (perhaps with accompanying sand and dirt). We will therefore evaluate the potential for the scanner bed to be scratched under regular use and, if necessary, investigate the use of either disposable films or a replacement scanner bed with a scratch-resistant coating.

Digital SLR Camera: In the event that using a flatbed scanner proves unworkable, our alternative is to use a high-quality digital SLR camera, mounted on a tripod and looking down at the fragment. The advantages are that the capture system does not come in contact with the fragment, so any potential issues with placing fragments on a sheet of glass are avoided.

We see three principal disadvantages to using a camera. First, the bottom of the fragment is covered in this scenario, making it impossible to capture 3-D shape simultaneously with color. From a technical standpoint, this complicates aligning the color data to the 3-D data. Second, we believe it will be harder to ensure that the camera is looking straight down on the fragment as opposed to at a slight angle. This too complicates aligning the data. Finally, it will be more difficult to control the lighting in order to obtain perfectly evenly-lit images.

2.3 Scanning Configuration

In order to obtain aligned color and 3-D data, and to provide an efficient scanning workflow, we must consider how the different pieces of acquisition equipment will be arranged. We will evaluate three possible configurations.

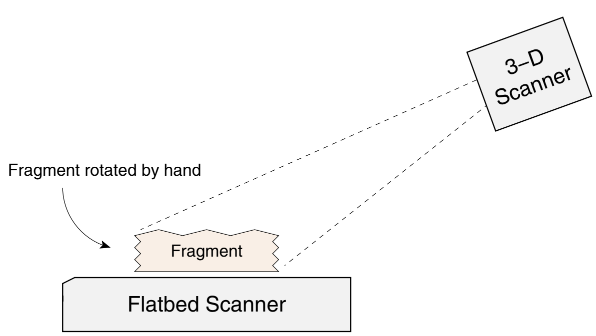

Scanning Configuration 1

The first of the proposed configurations uses a flatbed color scanner, on which the fragments are placed face-down. A 3-D scanner is mounted at an angle, to view the back of the fragment as well as one side. A single 3-D scan therefore captures only a part of the fragment, and multiple scans must be acquired and merged to form the complete model; typically, between 4 and 8 scans will be required. In this configuration, the fragment is moved by hand between scans.

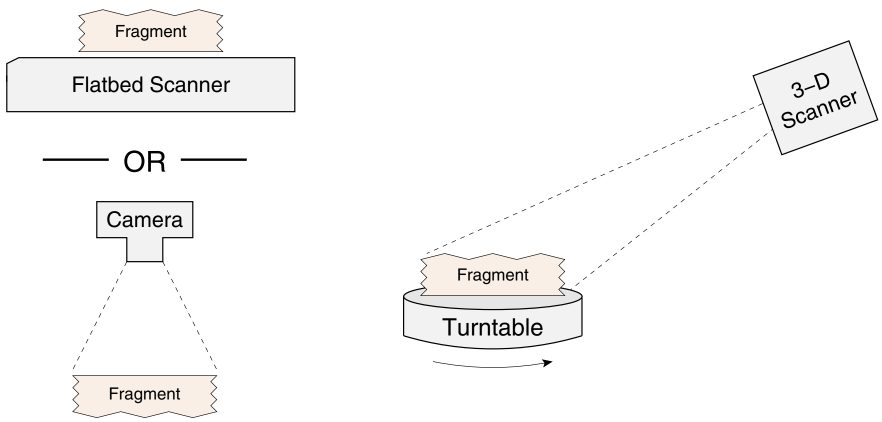

Scanning Configuration 2

The second proposed configuration removes the reliance on moving the fragment by hand. In this setup, an automated, computer-controlled "turntable" is used to hold the fragment and rotate it to different positions. This makes the scanning, as well as subsequent alignment of the separate 3-D scans, more efficient. However, because this configuration relies on a completely separate device to acquire color, it is more difficult to align the 3-D and color data.

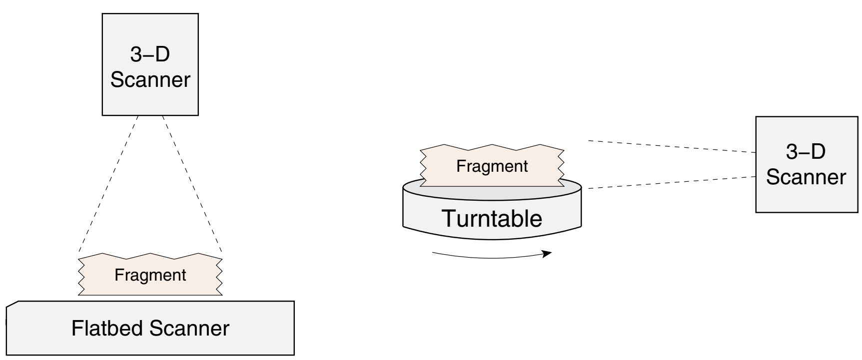

Scanning Configuration 3

The third configuration incorporates a 3-D scanner mounted in a calibrated way with respect to the color scanner. In this way, the 3-D shape of the back of a fragment is acquired at the same time as the color of the front. A second 3-D scanner, together with a turntable, are used to scan only the sides of the fragment. This still requires aligning two separate sources of 3-D data, but the problem is now one of 3-D to 3-D alignment, which is easier to solve than the 3-D to color alignment required in configuration 2.

Of the three proposed configurations, we believe that the third offers the potential for highest quality, efficiency, and flexibility, despite the additional equipment required. The final decision will be made based on our initial prototypes, and will be taken in cooperation with the archaeologists and conservators at Akrotiri. In the end, we anticipate that the acquisition system will be capable of acquiring complete color and 3-D information in a few minutes per fragment.

3 Fragment Matching

Currently, over 70% of the archaeological effort is devoted to assembling fresco fragments. We believe that computer matching algorithms can greatly reduce this effort. We will therefore provide a system that proposes matches to the archaeologist for confirmation, as well as a "virtual assembly table." In addition to saving time, digital assembly can also reduce the amount that fragments must be handled, which should improve their preservation. There are many cues on which matches may be computed, and we will explore all of them:

Location and Other Annotations. In many, if not most cases, adjacent fragments from the fresco are likely to have ended up in nearby places on the ground. Therefore, we will begin by considering information provided about where fragments were excavated, together with any additional available information about where on the wall the fragment is expected to be, the motif of the fragment, or digital marks "drawn" onto the scanned fragments. This will provide an initial grouping of similar fragments, reducing the number of potential matches that need be considered subsequently.

Thickness. Although thickness is not enough to guarantee a match, neighboring fragments should have the same thickness along their matching edges. Since this is a very simple criterion, it can be used to robustly prune away large numbers of implausible matches. This has been successfully used to guide matching in the Forma Urbis Romae project, and we expect it will be useful here too. Exactly how useful it will be will depend on how much the plaster varied in thickness originally, and how it has been affected by erosion over the centuries.

3-D Edge Profile. Manual matching relies heavily on finding fragments that "lock" together. Physically feeling this match is the most certain way of verifying its correctness. Since two non-matching edges can be expected not to have perfectly complementary, rough edges, this is also the most unique feature that can be used for computer matching. However, this sort of matching is stymied by edge erosion and cannot identify fragments that belong near each other but do not quite touch. Some work has been done in this direction, such as for the Forma Urbis Romae, but it remains a challenging problem that we will be researching further.

2-D Edge Contours. Less unique than the 3-D edge profile, but nevertheless a very important cue for matching, is the 2-D edge contour. Matching with this cue is similar to assembling a jigsaw puzzle in which all the pieces are blank. Substantial progress has already been made on this type of matching, and there is a significant body of literature on the subject. While carrying less information than the 3-D profile, the contour is also somewhat more robust to erosion.

Edge Color Matching. In addition to matching contours, color information from the fresco can be incorporated into the matching problem, leading to a more traditional jigsaw-assembly problem. Color can include both the color of painting and the discoloration of plaster. The Akrotiri frescoes use only a small number of pigments, so simple clustering techniques can robustly label pixels as "black," "red," or "yellow," leading to simpler color matching.

Pattern Continuation. Much of the fresco design is geometric, and researchers at NTUA have determined the mathematical equations governing these patterns. This makes it relatively easy to extrapolate designs well past the edge of a fragment. Then fragments that continue the pattern can be found. We expect this to be a particularly important technique for identifying nearby, but nonadjacent, fragments.

4 User Interface and Database

Fragments, once scanned, will be entered into a database and made available to software running on local or remote computers. This software application will:

- display all available information about each fragment. Options will be provided to display the color, thickness, edge shape, or any combination of these cues.
- allow users to move and rotate fragments on the screen. When two fragments are placed near each other, the quality of the match will be evaluated according to all of the metrics listed above, and the results presented to the user.
- permit users to add or edit annotations about any fragment.
- flexibly allow users to "group" and "ungroup" fragments, treating sets of matched fragments together.
- allow users to search for matches by simply clicking on a fragment. The search will be performed, and a ranked list of candidate matches will be presented for verification.
- record all alignment efforts that have been performed, including by multiple researchers working simultaneously on separate computers. The entire *history* of matching attempts will be recorded, allowing for an "undo" operation.

Once complete frescoes have been reconstructed, we will be able to construct a 3-D "virtual reality" model of an entire wall, room, or building, to be used as an aid to physical reconstruction or as an educational tool.

On the basis of past experience with systems that use computers to match fragments of archaeological artifacts, we believe that the proposed system will be able to quickly and automatically match the majority of adjoining fragments. In addition, we believe that the system may be able to find matches that people would not find, especially those involving small fragments with little or no color information. By permitting the majority of the fresco to be reconstructed within the computer, the proposed system will also reduce the amount of physical handling of fragments, reducing the wear to which these fragile objects are subjected.

References

[1] D. B. Cooper, A. Willis, S. Andrews, J. Baker, Y. Cao, D. Han, K. Kang, W. Kong, F. F. Leymarie, X. Orriols, S. Velipasalar, E. L. Vote, M. S. Joukowsky, B. B. Kimia, D. H. Laidlaw, and D. Mumford. Assembling virtual pots from 3d measurements of their fragments. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 241–254, 2001.

[2] C. Doumas. *The Wall-Paintings of Thera*. Thera Foundation, 1992.

[3] W. L. Friedrich, B. Kromer, M. Friedrich, J. Heinemeier, T. Pfeiffer, and S. Talamo. Santorini eruption radiocarbon dated to 1627–1600 B.C. *Science*, 312(5773):548, Apr. 2006.

[4] M. Kampel and R. Sablatnig. Profile-based pottery reconstruction. *Computer Vision and Pattern Recognition Workshop*, 2003.

[5] D. Koller and M. Levoy. Computer-aided reconstruction and new matches in the Forma Urbis Romae. In *Bullettino Della Commissione Archeologica Comunale di Roma*, 2006.

[6] D. Koller, J. Trimble, T. Najbjerg, N. Gelfand, and M. Levoy. Fragments of the city: Stanford's Digital Forma Urbis Romae project. In *Proceedings of the Third Williams Symposium on Classical Architecture, Journal of Roman Archaeology*, 2006.

[7] 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: 3D scanning of large statues. In *Proc. ACM SIGGRAPH*, pages 131–144, 2000.

[8] C. Papaodysseus, D. Fragoulis, M. Panagopoulos, T. Panagopoulos, P. Rousopoulos, M. Exarhos, and A. Skembris. Determination of the method of construction of 1650 B.C. wall paintings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, September 2006.

[9] C. Papaodysseus, T. Panagopoulos, M. Exarhos, C. Triantafyllou, and D. Fragoulis. Contour-shape based reconstruction of fragmented, 1600 B.C. wall paintings. *IEEE Transactions on Signal Processing*, 50(6):1277–1288, June 2002.

[10] Petra, the great temple excavations: Technology and excavations of the temple. `http://www.brown.edu/Departments/Anthropology/Petra/excavations/technology.html`.

[11] T. Vitale. Light levels used in modern flatbed scanners. *RLG DigiNews*, 2(5), 1998.

[12] A. Willis and D. Cooper. Estimating a-priori unknown 3D axially symmetric surfaces from noisy measurements of their fragments. In *Proceedings of the Third International Symposium on 3D Data Processing, Visualization, and Transmission*, 2006.