HIGH QUALITY 3D RECONSTRUCTION OF COMPLEX CULTURAL OBJECTS

M. Sormann[a], C. Zach[a], L. Zebedin[a], K. Karner[a]
[a] VRVis Research Center, Inffeldgasse 16/II, A–8010 Graz, Austria sormann@vrvis.at

KEY WORDS: orientation, segmentation, 3D reconstruction, texturing, photorealism

ABSTRACT

Nowadays, many cultural heritage applications require a 3D reconstruction of complex real world objects. Due to the complexity of real world objects, 3D reconstruction is often very time consuming and involves in general much manual effort. Therefore we report in this paper on a novel automatic 3D reconstruction approach to create virtual models of real world objects from a set of images with high redundancy. The power of our approach can be derived from the fact that we utilize the given redundancy to solve one of the main problems in photogrammetry, namely the correspondence problem. On the other hand we focus on minimizing the human interaction during the modeling process. The reconstruction pipeline starts with an image acquisition step, which consists of taking handheld images of the object with a calibrated digital consumer camera with short baselines. The rest of the workflow is represented by four consecutive steps: automatic orientation, multiple image segmentation, dense matching and multiview texturing. We demonstrate the accuracy, robustness and effectiveness of our proposed approach on several different real world objects.

1 INTRODUCTION

The preservation of complex cultural objects is an important application area of imagebased modeling. A full 3D reconstruction of a complex cultural object, like a statue, represents a permanent record of the real world object in their original position. The main interests for such a 3D reconstruction are the detection of erosion through environmental elements and the documentation of a restoration process. Another attractive application is the visualization of reconstructed 3D models in a virtual reality environment to make unique cultural objects accessible for a wide audience.

As introduced in ElHakim et. al. (6) the general requirements for the 3D reconstruction of real world objects are: high geometric accuracy, capturing all details, photorealism, full automation, low cost, portability and flexibility in applications. Our research aims to fulfil these requirements and to reconstruct a complex cultural object more accurate and convenient than current available computer vision techniques. Furthermore we focus on minimizing human interaction during the modeling process.

The images are captured directly on site with a high quality digital consumer camera using short baselines, resulting in high overlap. The key idea of our approach is to take advantage of the high redundancy to solve one the main problems in photogrammetry and computer vision, namely the correspondence problem. Essentially, the reconstruction pipeline can be separated into four main steps. The first step is represented by the automatic orientation procedure to obtain the relative orientation of each image pair. The next step consists of a multiple image segmentation task to separate the relevant parts of the scene belonging to the real world object from the background. The last two steps comprise the dense matching and the texturing of the 3D model, considering the visibility, viewing angle and the base of the projection.

The remainder of the paper is organized as follows. In section 2, an overview of related work concerning 3D reconstruction of complex cultural objects is provided. Our proposed 3D reconstruction approach is presented in section 3. Some results to demonstrate the usability of our approach are presented in section 4. The paper concludes with a short discussion and some aspects of future work.

2 RELATED WORK

The automatic reconstruction of complex cultural objects is still an active research field within the photogrammetric community. There are two classes of modeling techniques for acquiring 3D information of cultural sites. The first technique, known as rangebased modeling, is based on laser scanners. A very well known approach in this field is The Digital Michelangelo Project proposed by M. Levoy et.al. (14).

In this paper we focus on imagebased modeling, which represents the 3D reconstruction of real world objects from a dense set of photographs. A comparative evaluation of image-based and range-based methods can be found in ElHakim and Beraldin (5).

Imagebased modeling techniques utilize in general widely available hardware and a developed system can be normally used for a wide range of different objects and scenes. Furthermore such algorithms produce realistic models with an increasing level of automation. Over the past few years, a number of research teams have been addressed the use of these techniques to generate 3D models of complex cultural objects. One of the most popular approaches is the automated reconstruction of the Great Budha of Bamiyan, Afghanistan proposed by Gruen et.al. (9). The work is based on the use of three types of imagery in parallel: internet images, tourist images and metric images. The 3D reconstruction is performed with automatic and manual procedures utilizing the mentioned three types of data-sets. Another project of this research team is the 3D reconstruction and visualization of a Buddhistic Tower from Bayon Temple in Angkor, Cambodia proposed by Gruen et. al. (10).

The methods mentioned so far deal with several area-based matching techniques to recover the 3D information. A complete different image-based 3D modelling technique is shape form silhouette, which recovers the shape of the objects from their contours. A practical system to generate 3D models from its profiles was introduced by Wong and Cipolla (21). This approach uses only the silhouettes of a sculpture for both motion estimation and model reconstruction, and no corner detection nor matching is necessary.

3 3D RECONSTRUCTION OF COMPLEX CULTURAL OBJECTS

In this section we will report the incorporated steps of our novel automatic 3D reconstruction approach and present the underlying data structures and algorithms. As stated the input images are captured with a calibrated high quality digital consumer camera with a 11.4 megapixels CMOS sensor. The

image acquisition process consists of taking hand-held pictures with short baselines resulting in high overlap. The process of camera calibration, which is not the topic of the paper, is a well studied problem in photogrammetry and determines the internal parameters of a camera. Our method is based on work described by Heikkil¨a (12).
The remainder of the reconstruction pipeline works as follows. An automatic orientation procedure to obtain the relative orientation of each image pair is performed. A reliable calculation of the relative orientation is based on an accurate point of interest (POI) extraction followed by an affine invariant matching approach. The reconstruction of complex objects, like statues require a segmentation process to separate the relevant parts of the scene belonging to the statue from the background. This segmentation procedure consists of two consecutive steps. The first step involves some human assistance and is dedicated to identify the foreground or background in at least one image. Based on this initial segmentation an automatic segmentation propagation procedure is performed, which results in a segmentation of the whole image sequence. This 2D segmentation masks are important to reduce the outlier rate in the following dense matching procedure and to obtain a meaningful 3D reconstruction. Consequently, in our matching procedure we concentrate on combining the available silhouette information with an area based dense matching approach. Another important aspect to fulfil our 3D reconstruction requirements is a high quality texture of the 3D model, considering the visibility, viewing angle and the base of the projection. All these requirements are incorporated in our automatic texture generation method. The overall reconstruction pipeline is illustrated in Figure 1.

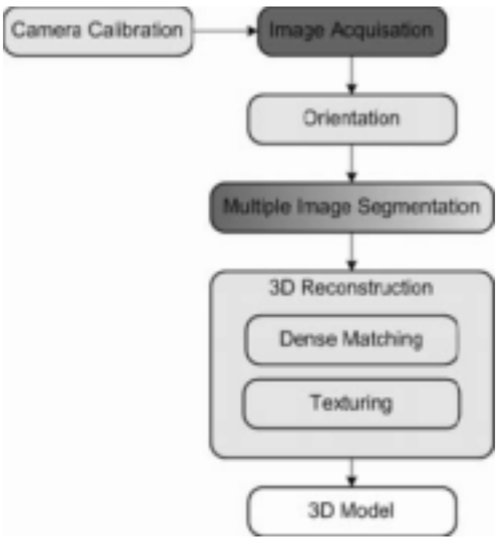

Figure 1: Overall reconstruction pipeline. The dark grey shaded tasks are interactive whereas the light grey shaded procedures are fully automatic.

All mentioned procedures of our 3D reconstruction pipeline will be explained in more detail in the following subsections.

3.1 Automatic Orientation

The first step in our 3D reconstruction approach consists of an accurate automatic orientation procedure to obtain the relative orientation of the whole set of images. One of the basic problems in the field of photogrammetry is the correspondence problem, which is to identify 2D points in two images that are projections of the same 3D point in the world. From corresponding points within the image sequences the relative orientation and the 3D positions of the corresponding points can be estimated.
Basically our orientation algorithm can be separated into four consecutive steps: extraction of tie points, matching of extracted tie points, estimation of unknown orientation parameters and refinement of obtained orientation parameters.
The tie point extraction is based on a state of the art point-of-interest detector proposed by Harris (11). Additionally each extracted tie point is provided with a feature vector, which consists of gradient information in a close neighbourhood. These feature vectors are utilized to detect corresponding points between image pairs. Another constraint is introduced by the well known epipolar geometry. Since the automatically matched correspondences can still contain a few outliers a robust estimation method called RANSAC (7) is initiated. In order to obtain the orientation of the whole image sequence, it is necessary to determine the scale factor. This task is accomplished by utilizing corresponding points in at last three images. The last step consists of a block bundle adjustment to refine the obtained orientation parameters. Figure 2 shows an oriented image sequence, illustrating the statue of St. Barbara, together with the obtained 3D tie points.

Figure 2: Illustration of an oriented image sequence together with the obtained 3D tie points. The image sequence consists of 47 images and approximately 5500 3D tie points which are shown as colored dots.

3.2 Multiple Image Segmentation

A detailed 3D reconstruction of complex real world objects requires many images, which in turn requires to segment many images which is a tedious and time consuming process. This section is dedicated to explain our multiple image segmentation task for minimizing the expenditure of time to achieve a robust and accurate foreground segmentation for 3D reconstruction. Such a foreground/background segmentation can simplify the correspondence problem dramatically and therefore, dense 3D reconstruction results of complex objects can be clearly improved. A more detailed description of the approach can be found in (19).
However, the multiple image segmentation procedure requires as input an initial segmentation of at least one image, which can be obtained by utilizing intelligent scissors introduced by Mortenson and Barrett (16) or other interactive segmentation techniques. Here we focus on the propagation of this initial segmentation through all images of the sequence. The propagation task itself is mainly based on a region based matching algorithm. Thus we segment the image into a certain number of regions. All these regions are classified into three different sets (foreground, background and uncertain regions), illustrated in Figure 3. The final contour can be extracted from these three sets. For dividing the image into regions we employ a mean-shift image segmentation proposed by Comaniciu and

Meer (3). Additionally, to improve the robustness of the propagation procedure, our algorithm requires the relative orientation of the images to be known.

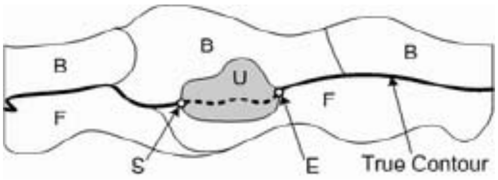

Figure 3: Set of regions including foreground regions (F), background regions (B) and uncertain regions (U) and the highlighted true contour. Furthermore an illustration of start point (S) and end point (E) to automatically apply intelligent scissors.

The workflow of our proposed segmentation method can be roughly seen as the composition of the following consecutive subtasks:

1. Extract an area of interest, which is called initial contour ring, with an inner boundary and an outer boundary.
2. Utilize information acquired from the contour ring and from the corresponding points to identify foreground and background regions with high confidence.
3. Perform a region based matching algorithm based on mean-shift information to separate the remaining regions in the contour ring in foreground regions, background regions and uncertain regions.
4. Extract true contour segments from adjacent foreground and background regions and utilize intelligent scissors to close uncertain contour segments.

This procedure is repeated until all images of the sequence are processed.

Essentially we label all mean-shift regions which are directly connected to the inner boundary of the ring to foreground and those connected to outer boundary as background. A similar procedure is performed for the information which is provided by the corresponding points. The core of our multiple image segmentation approach is represented by the extended region matching algorithm for the remaining regions. Therefore matching between two regions r_i of the previous image and r_j of the current image is assigned with a similarity measure $S_{i,j}$. The similarity measure $S_{i,j}$ is based on the mean-shift parameters and the known relative orientation. The final distance function for two regions is formulated as:

$$d(r_i, r_j)=\omega_1 * S_{LUV} + \omega_2 * S_{Epi} + \omega_3 * S_{Corr}$$

where S_{LUV} is represented by the LUV values of the mean-shift region, S_{Epi} encodes the distance of the epipolar line from region r_i to the center of gravity of region r_j and S_{Corr} is composed from the distance of the nearest corresponding point to region r_i respectively to region r_j . $\omega_1 ... \omega_3$ are weights to control the influence of the different similarity measures. By evaluating the introduced distance function for each remaining region against a user defined threshold, we distinguish between foreground regions, background regions and uncertain regions. For the uncertain regions the intelligent scissors algorithm is applied. If the segmentation result is not satisfactory, a user has the possibility to correct a miss-segmentation, by a manual assignment of critical mean-shift regions or by an assisted intelligent scissors algorithm. Figure 4 shows the intermediate results of the multiple image segmentation approach illustrated on the statue of St. Barbara.

3.3 Dense 3D Reconstruction

The set of images with known calibration and orientation is used to generate a 3D model of the object in a fully automated manner. Since the goal is the creation of proper 3D models, usual stereo methods generating only 2.5D heightfields are not sufficient. Alternatively, true multi-view reconstruction approaches such as space carving (13) or variational approaches (23) can be utilized. In order to increase robustness and speed of 3D reconstruction, we employ a two phase scheme to obtain a proper 3D model.

In the first pass a fast reconstruction method suitable for small-baseline settings is applied for every view. We utilize a plane-sweep approach (2, 22) to create the set of heightfields using up to 5 images simultaneously for matching (one key image in the middle and one or two neighboring reference images on each side). For each depth value, the reference images are projected onto the key image plane located at the given depth and a correlation measure with respect to the key image is calculated. Occlusion handling is addressed by the best half-sequence strategy. The set of slices filled with correlation values comprise a data structure similar to the disparity space image. A final matching algorithm (e.g. scanline optimization (18)) establishes the dense depth map from the disparity space image. Plane-sweep matching requires typically about 1.5 minutes for each reference image.

Figure 5(a)–(c) illustrate the results of plane-sweep matching. The range images are displayed as point clouds with color taken from the reference view. Obviously there are still mismatches and outliers in the generated heightfields. Removing these artefacts is postponed to the second phase of 3D reconstruction. Just merging the individual point clouds yields to the result shown in Figure 5(d), which is not suitable for further processing.

Silhouette information obtained from the previous segmentation step is used to restrict the heightfield to the foreground region of interest.

The next phase merges the individual heightfields into a final proper 3D model. We utilize a robust volumetric approach (4, 20) to combine the set of heightfields into a consistent 3D representation of the object. A prerequisite of this step is the specification of the 3D bounding box of the region to be reconstructed by the user.

Volumetric range image integration converts the range images into approximated volumetric distance fields and performs a robust averaging of these volumes. The final surface model can be extracted as the zero level of the averaged distance field. In order to reduce the memory requirements in this phase, we utilize a compact representation of the distance fields as described in (4). In spite of dealing with volumetric data structures, volumetric range image integration is surprisingly efficient. This step takes usually only a few minutes on standard hardware.

The extracted surface has usually a very high mesh complexity depending on the voxel resolution specified by the user. In order to reduce the mesh complexity any mesh simplification tool can be applied (e.g. (8)).

3.4 Texturing Complex 3D Objects

Since our goal is the generation of a photo-realistic model, an image (texture) has to be generated which can be mapped to the surface of the model. A perfect mapping between the triangulated surface and planar texture image should preserve angles and distances. Such a parameterization preserves obviously the areas of the triangles and is therefore optimally suited for texture mapping. Unfortunately general manifolds cannot be flattened without distortion. The deformation is roughly proportional to the amount of curvature the mesh exhibits.

The distortions caused by the parameterization can however be reduced by cutting the model into several smaller segments. This introduces discontinuities of the mapping along the seams. The visible effect of those errors can be alleviated by placing the cuts in regions of high curvature. In the final model such spots often involve brightness differences and therefore the additional discontinuity caused by the seam is masked by the change of illumination.

Once the number and shape of all regions has been determined, it is possible to find a mapping, which assigns each vertex a point in a plane. This function is often called a 2D-parameterization. Afterwards the different mapped pieces are packed into one image and textured. This is done by calculating the corresponding three dimensional point for every pixel which is projected into all available images. By applying an appropriate heuristic it is possible to estimate the real value and find a good approximation.

The algorithm employed in this work uses the approach of Levy (15) which is based on conformal maps to obtain the mapping of the three dimensional model to the two dimensional domain. The basic idea is to enforce the CauchyRiemann equation by minimizing it at the vertices of all triangles. The problem of combining different candidate pixels to approximate the real color property of a surface is discussed by Ofek et al. (17). They demonstrate that one valuable approach to this problem is to take the median of all possible values, thus eliminating highlights and reflection artefacts.

Bornik et al. (1)applied this method already successfully, but reconstructed textures only for planar (but not arbitrarily shaped) models. One additional feature our texturing process introduces is the noise and occlusion suppression quality. Considering the distance between the surface, the angle between surface normal and viewing direction of the camera and the general shape of the surface, one candidate color can be computed per image. The set of all resulting colors can then be further processed to obtain a good approximation. Once the visibility between the camera and the pixel is calculated, we can encounter three cases:

1. The point is not visible in this camera.
2. The point is visible in this camera, but was affected by some nonmodeled occlusion.
3. The point is visible in this camera.

In the first case the image can simply be discarded because it is obvious that it can not provide any information about the true pixel values. The last case also does not make any problems because it clearly has some valuable informations about the correct color we should use for the original pixel. The second case is the most difficult one to deal with because it can not automatically be discovered that there was an obstacle occluding the model. A very simple approach to solve this problem would be to assume that only very few such occlusions occurred (maybe this can be guaranteed by taking the images carefully) and calculate the average of all possible candidates. Thus the noise and the effect of outliers caused by obstacles would be reduced.

A more sophisticated approach combines all available information from the other images and create a heuristic that can be applied to calculate a robust approximation for the true texture. To make this possible we have to assume that there are at least some other images which have a valid visibility. Provided that we have some redundancy in the images and the point is visible in some of them, then we assume that only a minority of them have some obstacle in between which is not included in the model. We separate the candidates from the various input images into two categories in a next step. One category contains all gray-scale entries, the other the colored ones.

(a) (b) (c) (d) (e)

Figure 4: Intermediate results of the multiple image segmentation algorithm illustrating one image of the St. Barbara image sequence. **(b)** Closeup from **(a)** showing labelled foreground (green) and background (blue) regions in the contour ring after incorporating prior information. **(d)** Closeup from **(c)** illustrating labelled foreground (green), background (blue) and uncertain regions (red) after applying extended region matching. **(e)** Statue of St. Barbara with the final segmentation overlayed.

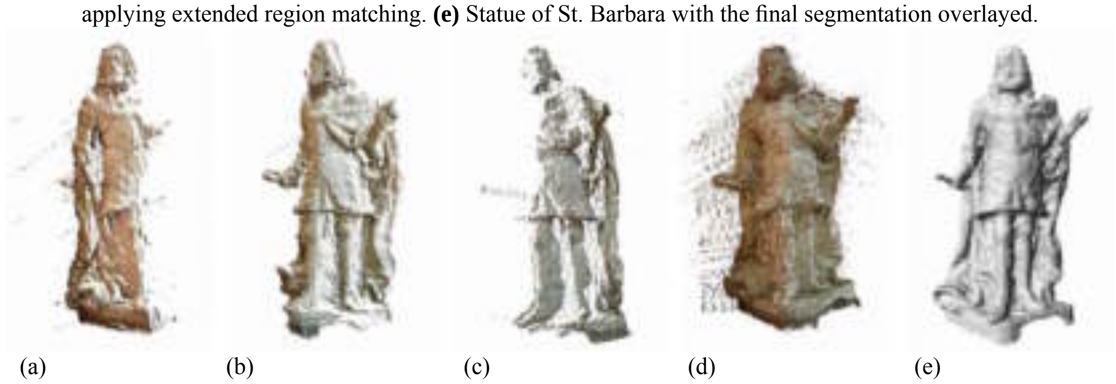

(a) (b) (c) (d) (e)

Figure 5: Three result heightfields of planesweep matching displayed as colored point set (a)–(c). Sole merging of the resulting point sets yield to the model displayed in (d). Robust volumetric integration of all generated heightfield results in the final model shown in (e).

This is done, because gray-scale entries are dealt with differently than colored ones. The distinction is made by converting the RGB color into the HSV color space. There the saturation is compared to some low threshold. This separation is also useful because colored outliers in otherwise grayscaled textures can easily be detected. Depending on which category has more members, an approximation is calculated based solely on the members of that category. This decision is repeated for the candidates of each texel.

4 RESULTS

Here we apply our approach to a variety of data-sets to demonstrate the usability and robustness of our method. In all of the experiments, the image resolution was 4064x2704 pixels. In order to reduce the mesh complexity a standard mesh simplification algorithm is performed, thus that all reported results consists of about 100.000 triangles.
In a first evaluation we used an image sequence consisting of 47 images of a statue of St. Barbara. The statue is 55cm tall with a diameter of approximately 13cm at the pedestal. Figure 6 illustrate our obtained reconstruction result of the statue of St. Barbara from different viewpoints.
Figure 7 shows a more complex data-set consisting of 45 images of the emperor Kaiser Karl VI. The statue of Emperor Karl VI. is
2.3 metres high and is located in the middle of the Great Hall of the Austrian National Library in Vienna.

5 CONCLUSION AND FUTURE WORK

A multi-stage approach for detailed and accurate 3D reconstruction of complex cultural objects was presented. Our approach takes advantage of the redundant scene information, which is typically provided by image sequences, captured with short baselines. Furthermore we focus on minimizing the human interaction during the modeling process. The reconstruction pipeline consists of four consecutive steps, which are known as automatic orientation, multiple image segmentation, dense matching and multiview texturing. Additionally, the presented method is demonstrated on different real world scenes to emphasize the feasibility of our implemented approach.
Though the results are very promising, there are several improvements that can be made to our approach. In order to improve the reconstruction quality it is necessary to deal with occluded parts of the object. Such occluded parts are usually leading to holes and introducing undesirable artefacts in the resulting models. Therefore we are currently working on a hole-filling strategy for reconstructed 3D surfaces. Moreover an efficient inpainting approach to enhance the texture quality have to be developed as well. Further we will concentrate on improving the performance of the developed algorithms.

ACKNOWLEDGMENTS

This work is partly funded by the VRVis Research Center, Graz and Vienna/Austria (http://www.vrvis.at). We would also like to thank the Vienna Science and Technology Fund (WWTF) for supporting our work in the the *Creative Histories The Josefsplatz Experience* project.

REFERENCES

BORNIK, A., KARNER, K., BAUER, J., LEBERL, F., AND H.MAYER. High-quality texture reconstruction from multiple views. *IEEE Transactions on Visualization and Computer Graphics 12*, 5 (October 2001), 263–276.

COLLINS, R. T. A space-sweep approach to true multi-image matching. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (San Francisco, USA, June 1996), pp. 358–363.

COMANICIU, D., AND MEER, P. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24*, 5 (May 2002), 603–619.

CURLESS, B., AND LEVOY, M. A volumetric method for building complex models from range images. In *ACM SIGGRAPH* (New York, USA, January 1996), pp. 303–312.

ELHAKIM, S., AND BERALDIN, J. Configuration analysis for sensor integration. In *Proceedings of SPIE* (Philadelphia, USA, October 1995), vol. 2, pp. 274–285.

ELHAKIM, S., AND BERALDIN, J. Detailed 3d reconstruction of monuments using multiple techniques. In *Workshop on Scanning for Cultural Heritage Recording* (Corfu, Greece, September 2002), vol. 1, pp. 58–64.

FISCHLER, M. A., AND BOLLES, R. C. Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Communication Association and Computing Machine 24*, 6 (June 1981), 381–395.

GARLAND, M., AND HECKBERT, P. S. Surface simplification using quadric error metrics. In *ACM SIGGRAPH* (Los Angeles, USA, August 1997), pp. 209–216.

GRUEN, A., REMONDINO, F., AND ZHANG, L. Image-based automated reconstruction of the great buddha of bamiyan, afghanistan. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Madison, USA, June 2003). CDROM.

GRUEN, A., ZHANG, L., AND VISNOVCOVA, J. Automatic reconstruction and visualization of a complex buddha tower of bayon, angkor, cambodia. In *International Conference on Optical 3D Measurement Techniques* (Vienna, Austria, October 2001), pp. 289–301.

HARRIS, C., AND STEPHENS, M. A combined corner and edge detector. In *4th Alvey Vision Conference* (Manchester, UK, September 1988), vol. 1, pp. 147–151.

HEIKKILA, J. Geometric camera calibration using circular con trol points. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22*, 10 (October 2000), 1066–1077.

KUTULAKOS, K., AND SEITZ, S. A theory of shape by space carving. *International Journal of Computer Vision 38*, 3 (July 2000), 198–216.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. The digital michelangelo project: 3d scanning of large statues. In *ACM SIGGRAPH* (New Orleans, USA, July 2000), vol. 1, pp. 131– 144.

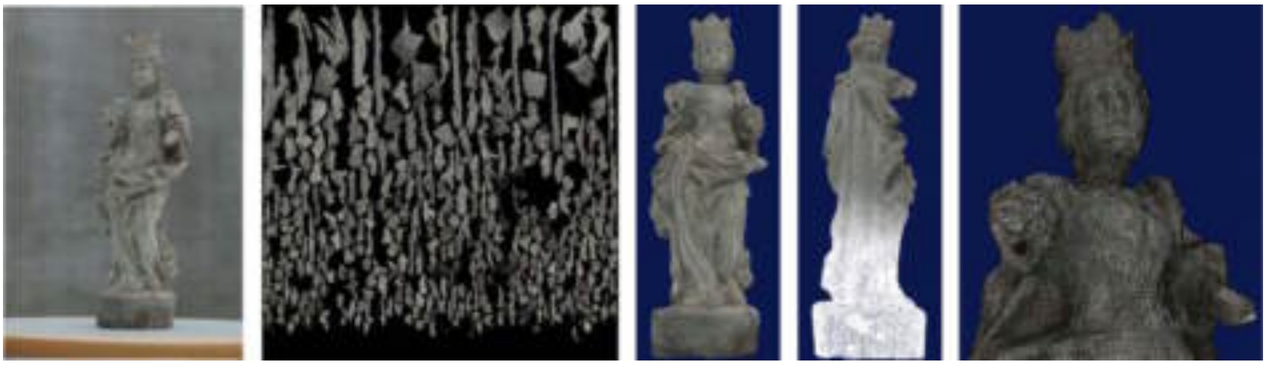

(a) (b) (c) (d) (e)

Figure 6: Dense 3D reconstruction of the statue of St. Barbara. The statue is 55cm tall with a diameter of approximately 13cm at the pedestal. **(a)** One original image of the statue of St. Barbara. **(b)** Generated Texture Map. **(c)-(d)** Different viewpoints of the reconstructed 3D model. **(e)** Close-up with overlayed wireframe to illustrate the geometric details.

(a) (b) (c) (d)

Figure 7: Illustration of all results achieved with our 3D reconstruction approach and demonstrated on the emperor Kaiser Karl VI located in Austrian National Library in Vienna. The statue of Emperor Karl VI. is 2.3 metres high and the dataset consists of 45 images. **(a)** One original image of the dataset. **(b)** Obtained camera positions and colored 3D tie points. **(c)-(d)** Two viewpoints of the reconstructed 3D model.

LEVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. Least squares conformal maps for automatic texture atlas generation. *ACM Transactions on Graphics 21*, 3 (July 2002), 362–371.

MORTENSON, E. N., AND BARRETT, W. Intelligent scissors for image composition. *Graphical Models and Image Processing 60*, 5 (September 1998), 349–384.

OFEK, E., SHILAT, E., RAPPOPORT, A., AND WERMAN, M. Multiresolution textures from image sequences. *Computer Vision, Graphics and Image Processing 17*, 2 (June 1997), 18–29.

SCHARSTEIN, D., AND SZELISKI, R. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision 47*, 13 (April 2002), 7–42.

SORMANN, M., ZACH, C., BAUER, J., KARNER, K., AND BISCHOF, H. Automatic foreground propagation in image sequences for 3d reconstruction. In *German Association for Pattern Recognition* (Vienna, Austria, August 2005).

WHEELER, M., SATO, Y., AND IKEUCHI, K. Consensus surfaces for modeling 3d objects from multiple range images. In *International Conference of Computer Vision* (Bombay, India, January 1998), pp. 917 – 924.

WONG, K.Y., AND CIPOLLA, R. Reconstruction of outdoor sculptures from silhouettes under approximate circular motion of an uncalibrated hand-held camera. *IEEE Transactions on Information and Systems 87*, 1 (January 2004), 27–33.

YANG, R., WELCH, G., AND BISHOP, G. Real-time consensus based scene reconstruction using commodity graphics hardware. In *Proceedings of Pacific Graphics* (Beijing, China, October 2002), pp. 358–363.

YEZZI, A., AND SOATTO, S. Stereoscopic segmentation. *International Journal of Computer Vision 53*, 1 (June 2003), 31–43.