

Preserving the Khmer Smile: Classifying and Restoring the Faces of Bayon

Min Lu¹, Bo Zheng¹, Jun Takamatsu² Ko Nishino³ and Katsushi Ikeuchi¹

¹Institute of Industrial Science, The University of Tokyo, Japan

²Graduate School of Information Science, Nara Institute of Science and Technology, Japan

³Department of Computer Science, Drexel University, USA

Abstract

The Bayon temple is known for its numerous massive stone faces with serene smiles, often referred to as the “Khmer Smile.” Many of these sculptures are, however, only partially preserved, making the desire to see the original appearance of Bayon faces elusive. We propose a novel method to restore these Bayon faces, which builds upon the matrix recovery theory. The method achieves accurate restoration by adopting a two-step shape recovery strategy. Rough restoration and clustering processes are first carried out using the entire database to group similar samples together. Then refined restoration using high resolution data is executed in each cluster to restore higher details while retaining the characteristics of each face. Experimental results demonstrate the effectiveness of our proposed method.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

1. Introduction

Historic cultural relics are often only partially preserved due to various causes, such as weather and vandalism. In order to avoid a great loss of cultural heritage, the issue of finding effective ways to restore the original appearance of these relics is important. When we have a group of similar relics and some of them are incomplete, can we restore them automatically using their geometric priors?

One example of such a problem is the famous facial sculptures in Bayon at Angkor. As one of the most important archaeological site in Cambodia, Bayon contains a multitude of massive stone faces on many towers, which are also the most recognizable and distinctive symbol of this temple. Staring out in all directions with serene smiles, these faces are often referred to as the “Khmer Smile.”

The temple itself, which was built in the late 12th or early 13th century, was carefully designed, with rich and beautiful decorations. By means of modern 3D scanning techniques, the entire building was digitally recorded and is known as the Bayon Digital Archival Project conducted by Ikeuchi et al. [IM07]. Images from arbitrary viewpoints can easily be

rendered using the scanned 3D data, such as those shown in the bottom row of Figure 1.



Figure 1: Images of Bayon, including an overview (left) and one face tower (right). The upper two are photographs and the bottom ones are rendered images using scanned 3D data.



Figure 2: Examples of Bayon faces. Many of them are incomplete, due to natural decay and damage from vandalism.

At the time of completion, the Bayon temple contained about 200 huge stone faces. Now only a certain number of them are partially preserved, as shown in Figure 2. Most damaged faces are still not restored, despite extensive efforts since the 20th century. In this paper, we reshape the Bayon faces to their former glory, using the incomplete data.

1.1. Related Work

Extensive studies of the Bayon faces have been conducted from the cultural and archaeological perspectives. A groundbreaking study on the classification of Bayon faces was done by the Japanese Government Team for Safeguarding Angkor (JSA) [Nak98]. Based on the observation and analysis from their experienced experts, those faces were roughly divided into three groups: *Deva*, *Devata* and *Asura*, meaning *god*, *goddess* and *devil* respectively. Typical representatives of these three groups are shown in Figure 3. A more objective classification study of Bayon faces was established by Kamakura et al. [KOT105], Using depth images generated from scanned 3D data.

The task of estimating the original Bayon face shapes relates to the 3D shape restoration problem. For the purpose of filling holes on given polygon meshes or point clouds, either mesh-based methods [Lie03] and volumetric approaches [DMGL02, NT03, Ju04] can be used to achieve smooth continuation by imposed localized geometric constraints. When large regions or structural information are lost, hole-filling methods will be no longer suitable. An alternative strategy using geometric priors can be employed. Usually patches with similar surface characteristics are selected from either the incomplete object itself [SACO04, BF08], or analogous candidate models [PMG*05, KS05]. Notice that all these methods handle a single sample at a time.

In case a group of similar shapes are present and a certain number of them are incomplete, an algorithm for simultaneously restoring all those shapes was developed by Lu et al. [LZT*10]. Using dense correspondences acquired from a shape matching scheme, the shape restoration problem was




Figure 3: Representative faces of the three categories in Bayon manually labeled by JSA [Nak98]. Left: *Deva* (No. 51S). Middle: *Devata* (No. 50E). Right: *Asura* (No. 35N).

formulated as a convex optimization and could be solved effectively. The pipeline of this method is illustrated in Figure 4. If there is more than one category among input samples, however, this basic algorithm should not directly be applied to the entire database. Otherwise it will bring the risk of restoring an individual using dissimilar samples from different categories, which is unreasonable. In this case, shape recovery processing should be carried out in a relatively small scope, where only samples with high similarity values are chosen, in order to retain the characteristics of each face and get accurate restoration results. In addition, this method does not distinguish data loss and damage, which will lead to inaccurate output under certain conditions.




Figure 4: The pipeline of the shape restoration method [LZT*10]. Input 3D shapes are first aligned and dense correspondences are acquired using a surface matching scheme. By stacking coordinates of these corresponding points, input shapes are represented as fixed-length vectors and then a matrix recovery procedure is used to accomplish the restoration.

1.2. Algorithm Overview

In this paper, we propose a novel comprehensive method to repair Bayon faces. We leverage the fact that these faces are exactly a group of similar shapes and base our method on the restoration algorithm proposed by Lu et al. [LZT*10].

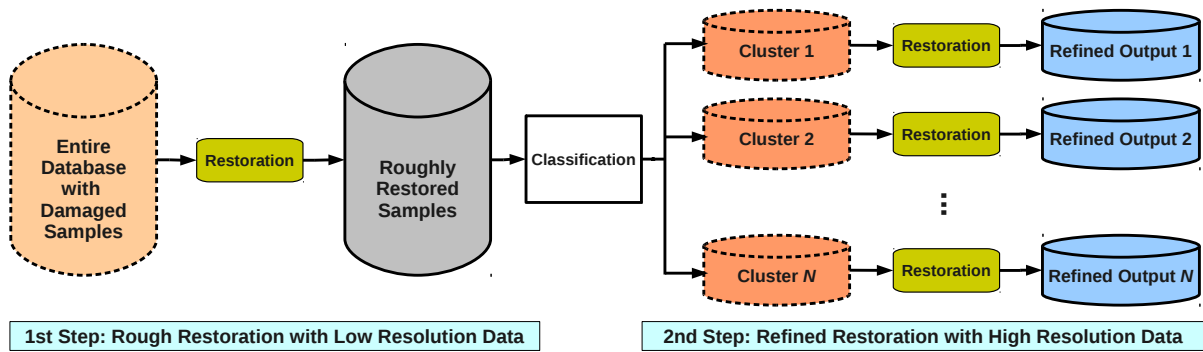


Figure 5: The workflow of our proposed two-step restoration method. Notice that the same restoration algorithm, which is explained in Section 3.2, is used in both steps.

Considering that these sculptures belong to several different categories, a two-step restoration strategy is developed, as shown in Figure 5. First, in order to obtain more accurate classification results, where similar faces are grouped together, a rough shape recovery using all samples is carried out to temporarily handle data incompleteness. Second, a refined restoration using the same algorithm but with high resolution data is executed in each respective cluster, which results in detailed restorations that retain the characteristics of each face. Several advantages are brought by the proposed method, including higher restoration accuracy and reduced computation.

The remainder of this paper is structured as follows: First, the cluster analysis of Bayon faces is introduced in Section 2. Then the restoration method, both the original and our improved versions, is explained in Section 3. Experimental results are shown in Section 4 and the last section concludes the paper.

2. Classification

2.1. Pre-processing

The database of Bayon faces coordinated by Ikeuchi et al. [IM07] contains registered 3D scans of 161 faces. Before starting further analysis, these 3D shapes have to be pre-processed to fit the requirements of our proposed algorithm.

By adopting a surface matching scheme [LZT*10], dense corresponding points are generated. Notice that the outer parts, such as ears and headdresses, are cropped. Only the inner part, the face itself, is left. Then we stack the (x, y, z) coordinates of these points, turning each face into a fixed-length vector to represent its geometric shape.

2.2. Hierarchical Clustering

To achieve detailed restoration results, we must first group the faces into similar shapes. This is a typical task of cluster analysis, with several classical solutions, such as K-means [Bis06]. Notice that for our problem, although we all believe these faces should be classified, it is difficult to determine the exact number of how many clusters they should be divided into. Here the hierarchical clustering method [DHS01] is chosen. Instead of generating a “flat” data description, like K-means, this method will lead to a hierarchical representation, where the relationship among all samples is clear at a glance and clusters can be chosen more flexibly.

Given pairwise similarities of all samples, a hierarchy is built from the individual faces by progressively merging clusters. Ward’s criterion [JHW63] is used as the linkage criteria which determines the distance between pairs of clusters. The result is represented in a corresponding tree, called a dendrogram, as shown in Figure 6. Note that the similarity values, which correspond to the height in the dendrogram, can be used to help determine whether groupings are natural or forced. For instance, if an unusually large gap between the similarity values of two neighboring levels is found, one can argue that here is a natural boundary to divide the dendrogram.

In the practice of refined restoration, an additional benefit is that for a certain face, which is one leaf node in the dendrogram, the scope of the cluster used for recovery can be gradually enlarged, until a satisfactory result appears. This bottom-up strategy also enables efficient computation on limited memory, as only a part of the database is required in the computation.

After the pre-processing explained in Section 2.1, the similarity evaluation between two Bayon faces becomes just measuring the distance between two points in a vector space for which we use the Euclidean norm. The raw 3D measure-




Figure 6: The hierarchy of Bayon faces. Based on the global rough restoration results using low resolution data.

ments, however, should not be used to measure similarities of these faces as they may be incomplete. Instead, a rough shape restoration procedure is carried out beforehand, using all but relatively low resolution data. The details of the restoration algorithm is explained next.

3. Restoration

3.1. Shape Restoration via Matrix Recovery

For a certain missing or damaged part of one Bayon face, a relatively intact patch corresponding to the same area can always be found from at least another sample in the entire database. To leverage this fact, we base our method on that of Lu et al. [LZT*10] which we briefly explain for completeness.

The aim is to get rid of corruption as much as possible and recover undamaged shapes. Given a sufficiently large number of similar shapes, let $\{\mathbf{s}_i^0\}_{i=1}^n$ and $\{\mathbf{s}_i\}_{i=1}^n$ denote their original shapes when they were just completed and the corresponding observations, i.e., raw measured data in real world respectively.

An assumption, that $\{\mathbf{s}_i^0\}_{i=1}^n$ are linearly correlated, is required as a premise. Let $\text{span}(\mathbf{s}_1^0, \dots, \mathbf{s}_n^0)$ denote the linear span of these shapes:

$$\text{span}(\mathbf{s}_1^0, \dots, \mathbf{s}_n^0) = \left\{ \sum_{i=1}^n \alpha_i \mathbf{s}_i^0 \mid \alpha_1, \dots, \alpha_n \in \mathbb{R} \right\}, \quad (1)$$

which represents the intersection of all subspaces containing this set. The assumption of linearity means the dimension of this linear span should be much smaller compared with the superposition of n times the value of each sample. In other words, if vectors $\{\mathbf{s}_i^0\}_{i=1}^n$ form the columns of a matrix, denoted as A , then the rank of this matrix should be much less than the number of samples n , which means A is approximately a low-rank matrix.

Similarly, an observation data matrix D is formed by vectors $\{\mathbf{s}_i\}_{i=1}^n$. The difference between A and D corresponds to the corruption, which is denoted as matrix E . The observation can then be decomposed as

$$A + E = D. \quad (2)$$

Usually defect area is much smaller comparing with the complete parts, which means error matrix E should be sparse and most of its entries are zero. Considering that matrix A is approximately low-rank, as analyzed earlier, the task of restoring this group of similar shapes can be formulated as a low-rank matrix recovery problem [CLMW11]:

$$\min_{A, E} \|A\|_* + \lambda \|E\|_1, \quad \text{s.t. } A + E = D, \quad (3)$$

where $\|A\|_*$ is the nuclear norm of matrix A , defined as the sum of its singular values: $\|A\|_* \doteq \sum_i \sigma_i(A)$. $\|E\|_1$ refers to the L_1 norm of matrix E . This problem can be efficiently solved by convex optimization, for example, using Augmented Lagrange Multiplier (ALM) algorithm [LCM10].

3.2. Shape Restoration via Matrix Recovery and Completion

The basic recovery method treats missing and damaged parts the same, leading to inaccurate restorations when directly applied to the Bayon faces. For instance, if the missing parts of one sample are considerably large, while the total number of samples are limited, which means the size of data matrices is relatively small, the corruption matrix E would be no longer sparse if the missing parts are not excluded. Inaccurate results may be generated in this case. In order to improve the restoration method to avoid this problem, cases of data loss and corruption have to be distinguished, and treated differently.

Suppose the size of matrix D in Eq. 3 is $m \times n$. Let Ω denote the positions of observed entries of D . $\Omega \subseteq [m] \times [n]$ and it is called the support set of D . Its complementary set Ω^c corresponds to those missing entries, which have been localized already after the surface matching procedure in preprocessing. Let P_Ω be an orthogonal projection supported on Ω , defined as

$$P_\Omega(X) = \begin{cases} X_{ij}, & (i, j) \in \Omega, \\ 0, & (i, j) \notin \Omega, \end{cases} \quad (4)$$

where $X_{m \times n}$ is an arbitrary matrix. This projection can be used as a mask to exclude missing entries of matrix E in

Eq. 3, making the rest of this difference matrix E sparse. The former optimization Eq. 3 is then modified as

$$\min_{A,E} \|A\|_* + \lambda \|P_{\Omega}(E)\|_1, \quad s.t. \quad A + E = D, \quad (5)$$

which is still a convex optimization and can be solved by the ALM method as well.

Note that we adopt this restoration method in both rough and refined shape recovery steps using data with different resolution, and the scopes of the faces within the database are different, too. Due to the cropping operation, this method handles the facial part only, without considering outer parts such as ears and headdresses. In the end, restored outputs are merged together with their initial shapes in order to obtain better appearance.

4. Experiments

4.1. Results of Revised Restoration Method

In Section 3.2, an improved version of the original shape restoration method is presented. The effectiveness of the modification can be seen in the following experimental results.



Figure 7: Examples of synthesized data. Upper left: the original shape. The rest: synthesized samples, generated by cutting down a certain facial area randomly, as illustrated in Figure 7. Holes are filled using simple smooth surface.

A classification test was designed to check the correctness of our proposed method. First, we built a small data set with known ground truth of grouping information. For each chosen face, five extra samples were synthesized by cutting down a certain area randomly, as illustrated in Figure 7. Holes were filled using simple smooth surface. Six well preserved faces in total were chosen from the entire database. This results in a data set containing 36 samples, on which both the algorithm [LZT*10] and our proposed method were carried out, aiming to restore those incomplete samples back to their origins.




Figure 8: An experiment on a synthesized data set for testing the correctness of the novel shape restoration method. Left: original shapes. Middle left: synthesized inputs, with nearly half of each face cropped. The missing part of each face corresponds to a plane on the surface of its bounding box. Middle right: corresponding restored outputs of the method [LZT*10]. Right: restoration results of the improved method we proposed.

Based on the procedure how this data set was built, six categories, each containing one original shape and five corresponding synthesized samples, should be found when choosing a proper scale for clustering the restored outputs, if the restoration algorithm works correctly. Notice that here the restoration was adopted only once using all 36 samples, because we just intended to show the correctness of our novel restoration method, not aiming to obtain the refined restoration results.

Part of the restored outputs are shown in Figure 8. Compared with the outputs of the former method, more accurate results were achieved by our revised algorithm. The outcomes of the former algorithm, shown in the third column of Figure 8, look very similar with each other and led to an incorrect clustering result in our experiment.

4.2. Results of the Two-Step Restoration

We next show the results of applying the proposed method on the entire Bayon face database, including 161 samples in total.

First, after the rough restoration and classification procedures, similar samples were grouped together to generate a hierarchy shown in Figure 6. Here we choose one cluster of the whole dendrogram, which is marked by a red box in that figure, to demonstrate the restored results of our method. The

enlarged image of the hierarchy corresponding to this cluster is shown in Figure 9.

There are 12 samples in the chosen cluster, where several samples are partially damaged, as shown in Figure 10. For example, the third face in the bottom row, which is numbered as No. 49E, has a broken nose. After applying the refined restoration process to this group, the missing nose part of this sample is recovered, as shown in Figure 11. From the comparison from both frontal and side views, as well as the enlarged partial details, we can clearly see that the broken nose and mouth are well restored, without changing other undamaged areas. Another two restoration examples are presented in Figure 12. Notice that only the central facial part was taken into consideration during the whole processing and the cropped outer parts were not modified.




Figure 9: Enlarged image of the hierarchy corresponding to the chosen cluster in Figure 6.

Two examples of the two-step restoration outputs are shown in Figure 13. The refined outputs benefited from higher resolution data and more details were kept than the rough restorations. Note that the rough output of the second example changed slightly and became not so that similar as its input face. This was caused by the side effect of utilizing the entire database for restoration at the same time, that dissimilar samples may be unreasonably used. This demonstrates the rationality of our proposed method, where refined restoration is carried out using similar samples only. More restored results can be found in Figure 14, including various examples of damaged faces.

4.3. Discussion

True quantitative evaluation of these results cannot be attained, as the original shapes of these sculptures when they were just completed, i.e. the ground truth are not known. In spite of this, from the subjective judgment of us human observers, the restoration method we proposed clearly shows



Figure 10: The 12 samples that belong to the chosen cluster in Figure 6.



Figure 11: One experimental result of our proposed restoration method, face No. 49E, with comparison from a side view and enlarged partial details. Upper: measured shape. Bottom: restored output.

its feasibility and effectiveness in the experimental results, which is meaningful for the further archaeological research of Bayon. Besides, the proposed method is actually designed for a group of similar shapes, not specific to the Bayon data set. It can be applied to other problems encountered in archaeology and cultural heritage as well, such as to restore and classify a group of human skull fossils or similar bronze mirrors.

5. Summary

In this paper, we focused on the problem of restoring Bayon facial sculptures, estimating the original shapes as they were



Figure 12: Two more examples of restored results, with comparison from a side view and enlarged partial details.




Figure 13: Two examples of the two-step restoration results. Each group shows the inputs (top) and outputs (bottom) of both rough (left) and refined (right) restoration procedures. The refined outputs benefited from higher resolution data and more details were kept. Note that the rough output of the second example is not so that similar as its input face, which was caused by the side effect of utilizing the entire database for restoration at the same time, that dissimilar samples may be unreasonably used.

constructed. In order to achieve this task, we propose a novel comprehensive method with a two-step restoration strategy. First, a rough shape recovery using all samples is carried out to temporarily handle data incompleteness. This is to make the classification result more accurate. A hierarchical clustering using the restored outcomes is adopted to group similar samples together. Then a refined restoration using the same shape recovery algorithm but with high resolution data is executed in each respective cluster. This results in detailed restorations and the characteristics of each face are retained as well. A revised shape restoration method based on the low-rank matrix recovery theory was also derived. Experimental results clearly show the feasibility and effectiveness of the proposed method.

Additional development will include estimating data confidence values for input shapes before restoration, so that only those areas with low confidence values, which means they are more likely to be damaged, will be focused on during restoration. Also we want to refine the current two-step restoration strategy and develop it into an iterative algorithm where the procedures of classification and restoration are adopted alternately, in order to obtain more accurate results. More experiments using data other than Bayon faces will also be attempted.

Acknowledgments

This work was supported in part by Ministry of Education, Culture, Sports, Science and Technology under Digital Museum Project to KI, and National Science Foundation CAREER Award IIS-0746717 and IIS-080367 to KN.

References

- [BF08] BRECKON T. P., FISHER R. B.: Three-dimensional surface relief completion via nonparametric techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30 (2008), 2249–2255. 2
- [Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006. 3
- [CLMW11] CANDÈS E. J., LI X., MA Y., WRIGHT J.: Robust principal component analysis? *Journal of the ACM* 58 (2011), 11:1–11:37. 4
- [DHS01] DUDA R. O., HART P. E., STORK D. G.: *Pattern Classification*, 2 ed. Wiley-Interscience, 2001. 3
- [DMGL02] DAVIS J., MARSCHNER S., GARR M., LEVOY M.: Filling holes in complex surfaces using volumetric diffusion. In *The First International Symposium on 3D Data Processing Visualization and Transmission* (2002), pp. 428 – 861. 2
- [IM07] IKEUCHI K., MIYAZAKI D. (Eds.): *Digitally Archiving Cultural Objects*. Springer-Verlag, 2007. 1, 3
- [JHW63] JOE H. WARD J.: Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association* 58, 301 (1963), 236–244. 3
- [Ju04] JU T.: Robust repair of polygonal models. *ACM Transactions on Graphics* 23, 3 (2004), 888–895. 2



Figure 14: More restoration results. Top: original shapes. Upper middle: high resolution data used as input, containing about 100,000 points each. Lower middle: the refined restoration results. At most 36 samples from a same cluster were used for generate one certain output. Bottom: the final restoration results after merging. Here meshes instead of point clouds are used in the two middle rows for better visualization.

- [KOTI05] KAMAKURA M., OISHI T., TAKAMATSU J., IKEUCHI K.: Classification of Bayon faces using 3D models. In *The 11th International Conference on Virtual Systems and Multimedia (VSMM' 05)* (2005). 2
- [KS05] KRAEVOY V., SHEFFER A.: Template-based mesh completion. In *Proceedings of the third Eurographics symposium on Geometry processing (SGP' 05)* (Aire-la-Ville, Switzerland, 2005), Eurographics Association. 2
- [LCM10] LIN Z., CHEN M., MA Y.: The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. *ArXiv e-prints* (2010). 4
- [Lie03] LIEPA P.: Filling holes in meshes. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing (SGP' 03)* (Aire-la-Ville, Switzerland, 2003), Eurographics Association, pp. 200–205. 2
- [LZT*10] LU M., ZHENG B., TAKAMATSU J., NISHINO K., IKEUCHI K.: 3D shape restoration via matrix recovery. In *The 10th Asian Conference on Computer Vision Workshop on e-Heritage* (2010). 2, 3, 4, 5
- [Nak98] NAKAGAWA T. (Ed.): *Annual report on the technical survey of Angkor Monument*. Japan International Cooperation Center, 1998. 2
- [NT03] NOORUDDIN F. S., TURK G.: Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics* 9, 2 (2003), 191–205. 2
- [PMG*05] PAULY M., MITRA N. J., GIESEN J., GROSS M., GUIBAS L. J.: Example-based 3D scan completion. In *Proceedings of the third Eurographics symposium on Geometry processing (SGP' 05)* (Aire-la-Ville, Switzerland, 2005), Eurographics Association. 2
- [SACO04] SHARF A., ALEXA M., COHEN-OR D.: Context-based surface completion. *ACM Transactions on Graphics* 23 (2004), 878–887. 2