

Projecting Restorations in Real-Time for Real-World Objects

Alvin J. Law, Daniel G. Aliaga, Yu Hong Yeung, Purdue University, U.S.A.

Richard McCoy, Indianapolis Museum of Art, U.S.A.

Amy McKune, Eiteljorg Museum of American Indians and Western Art, U.S.A

Larry Zimmerman, Indiana University-Purdue University Indianapolis and Eiteljorg Museum, U.S.A.

Abstract

In this paper we introduce a virtual restoration system that projects three-dimensional visualizations of plausible restoration states on cultural objects. Unlike typical 3-D restorations that restore an object on a video screen, our restoration approach allows the viewer to enjoy the phenomenological experience of viewing an object in person and considering a variety of possible restoration states. Using this system, users now have the ability to visually explore and compare different presentation states (restorations) without physically altering an object. These restorations are generated with minimal user interaction and employ a controlled amount of light. They are presented instantaneously and without the aid of a personal viewing apparatus. Further, multiple viewers can observe a restored object at one time. This ability to quickly suggest and compare multiple restoration states has the potential to greatly benefit the public, museum professionals, and historians.

Keywords: Virtual Restoration, Three-Dimensional Projection, Cultural Property

1. Introduction

In this paper we introduce a means to create and view high-quality restorations of cultural property. Using off-the-shelf hardware we have developed a fundamentally different way to visually explore and compare multiple presentation states (restorations) on objects without physical alteration. On a single object we can suggest an original, as-made appearance to an appearance that suggests or completes missing designs or forms. Because these restorations do not physically modify the object and are archival, our approach also allows for repeatable and flexible restorations.

The key advantage to virtual restoration and our approach is that it circumvents the potential for ethical dilemmas caused by physical restoration. Our virtual restorations do not require the physical alteration of the object and are generated with adjustable light sources. They are produced immediately and are completely reversible.

To accomplish a restoration we start by capturing a digital image of an object using a digital camera as the only input device. This image is then fed into a computer program which calculates the position, color, and intensity of light needed to project back on the object. Finally, a new image is projected onto the object from several digital projectors. This semi-automatic restoration process provides multiple viewers with results that can be viewed immediately and experienced firsthand without the aid of any viewing devices. Each restoration is produced within minutes with some interaction from the viewer. Additionally, our system has the option to provide virtual illumination on the object to highlight certain portions or surface details.

In the sections that follow, we discuss some prior work on the virtual restoration of cultural property and then provide a detailed overview of the steps required to generate our virtual restorations. We demonstrate and examine our system on a variety of objects and conclude with a discussion of its limitations and our future plans. For more complete information on our restoration process, we refer you to our technical paper (Aliaga et al. 2008).

2. Prior Work in Virtual Restoration

Some recent projects have digitized important cultural property such as the Michelangelo (Levoy et al., 2000) and the *Forma Urbis Romae* (Koller et al., 2006). Other digitization projects related to archaeology and restoration include the work by Kampel and Sablatnig (2003), Nautyal et al.(2003), Skousen (2004), and Viti (2003). In these examples the digitized objects are stored as computer models that can be viewed, rendered, and modified as desired. While these projects were successful in imaging these objects and suggesting possible appearances, the 3-D model remains digitized and trapped inside the video screen. Our process of virtual restoration is unique because it brings the virtual restoration to the object.

On the technical side, Raskar et al. (2001) presented a projector-based system that projects light onto custom-built objects. The objects were constructed so as to imitate the general form of a desired physical structure, and the projectors provided additional visual details. As opposed to our colored (and deteriorated) objects, their projection surfaces were white and smooth thus only requiring simple photometric calibration.

3. An Improved Virtual Restoration System

Our virtual restoration process starts by placing an object in front of the camera (Canon Digital Rebel XT_i 10MP) and projectors (Optoma EP910 1400 x 1050 DLP). Using our restoration program, images are captured of a grid-pattern that is projected on the object. These images are used to acquire a geometric model of the object. With this model we can then estimate the internal and external parameters that the projectors will need to project the virtual restoration. Next, the restoration program uses this geometric model and a diffusely-lit photograph of the object to produce a synthetic restoration of the object.



Fig. 1: Our virtual restoration system consists of self-calibrating, off-the-shelf hardware.

The user adjusts this synthetic restoration by specifying the amount of smoothness of the desired restoration, and then calculates the light-transport matrix and radiometric calibration. Finally, a compensation image for each projector is interactively computed using a desired set of restoration parameters and maximum light-intensity. These images are projected back onto the objects giving it the appearance of having been restored.

Our acquisition approach is based on Aliaga and Xu (2008), which estimates the geometric positions of a dense set of point samples, their corresponding surface normals, and the positions and focal lengths of the projectors. The result is a dense and accurate model containing position and normal data of the target object. Since the system is self-calibrating, it does not require the user to perform separate model acquisition and calibrations of the camera and projector.

Much of this process is fully automated and takes about 1-2 hours to complete per object, with half of the time spent acquiring images.

4. The Restoration Process

Our virtual restoration method uses an iterative energy minimization process. Figures 2a and 2b show a detail of a ceramic vessel in the Casas Grandes style with a few common condition issues (faded pigmentation, a minor chip, losses, and staining). Our goal is to arrive at Figure 2c, a plausible restoration of the object, and Figure 2d, which shows Figure 2a virtually restored. The end result will provide a representation of how the object might have looked shortly after it was made.

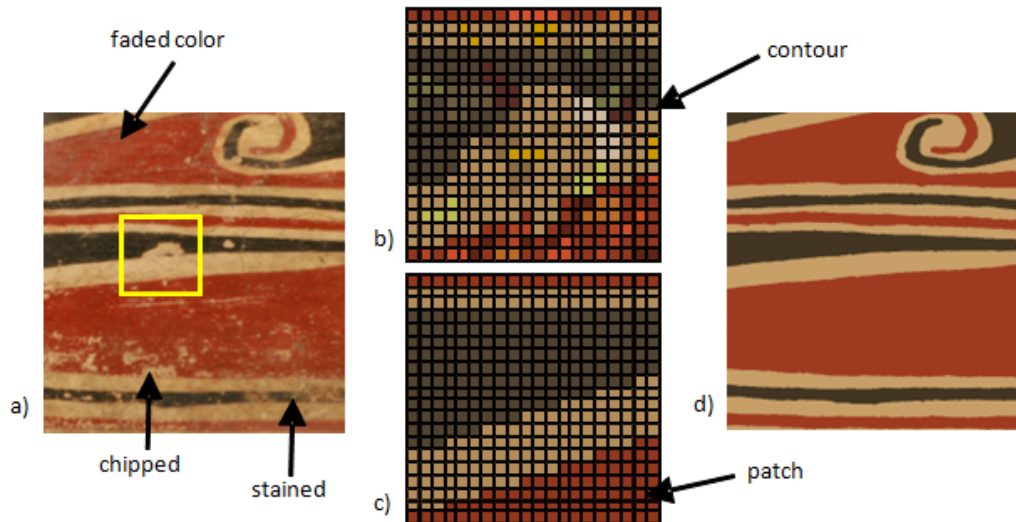


Fig. 2: Motivation for our restoration procedure. a) Detail (photograph) of a ceramic vessel in the Casas Grandes style. b) The synthetic equivalent of a pixel-subset of the image before restoration. c) The same synthetic region after our restoration. d) The final restored image of a).

Our system is able to work at any scale so as to provide both minor and major alterations to an object’s appearance. The user interactively controls the smoothness of the restoration as well as the degree that it deviates from the initial appearance.

To begin the restoration process, our color classification algorithm first assigns each pixel in the photograph of the object to one of a small set of colors. Then, the user selects a region of the photograph for restoration and assigns the region a background color. The size of the region depends on the object, and the total number of regions can range from one region to numerous regions. Within a region, each contiguous group of pixels of the same color forms a patch. The energy minimization alters the contours of the patches so as to reduce the value of a restoration energy term. For each patch in the region, the pixels of the patch’s contour are placed into a priority queue. The pixel with the largest energy term is processed first until either no significant change in energy occurs or no energy term exceeds a provided maximum. This process, combined with simple hole filling and noise removal, yield a restoration of the observed patterned object. Lastly, the computed image is used by the visual compensation method to calculate images for each projector.

Color classification

Using a photograph of the object as the input, our classification system places each pixel into one of a small set of colors by using a calibrated k-means clustering method. The photograph to be used for restoration is captured using indirect and diffused illumination so as to reduce the effects of shadows and highlights. Then, the user selects the number of representative colors n that are used by the object pattern and selects example pixels for each color from all over the object. Since most objects have few colors, this process typically takes only a few seconds and provides calibration pixels for improving a k-means clustering approach. Our classification scheme is capable of supporting larger numbers of discrete colors with additional pixel selection. However, as the number of colors increases, the reliability of color classification decreases.

Energy minimization

Our program contains an energy minimization routine which alters the contours of contiguous patches of pixels classified to the same color in order to

- smooth changes in the contour shape of each patch,
- smooth changes in the distance between the closest contour pixels of neighboring patches, and
- perform compensation-compliant contour changes.

Smooth contours generally yield pleasing results and can be assumed to be the case for hand-painted patterns. Encouraging smooth changes in the distance between pairs of closest contour pixels ensures smooth structural changes in the colored patterns. For example, if the opposing portions of two contours are irregular in shape but of roughly similar distances, then

the criteria will steer the contour pixels towards the spatial relationship of two parallel lines. However, if the contours are irregular and the closest pixel distance varies significantly and gradually from one side of the contour to the other, then the criteria will produce the desired straight line contours but also maintain the general spatial relationship between the contours. Further, since our goal is to ultimately restore the appearance of the object, ensuring the restoration is compensation compliant assures that radiometric compensation can be done well.

User input

During image restoration, the user provides input to the system to guide large structural changes in the observed patterns due mostly to large losses of decoration or severe color fading. Features which were originally part of the object may be too abraded to detect. In these cases, the user simply has to provide a rough sketch of the missing features on the image using a virtual paint-brush tool. These rough sketches do not have to be exact but only provide the proper connectivity of patches (e.g., if two patches are separated by a large area of missing paint, the user only needs to draw a colored line to connect the two patches). The later stages of restoration will inflate and/or deflate the shapes accordingly to satisfy the aforementioned restoration criteria.

This input by the user is the primary work required to generate a plausible restoration. Once this input is provided, the bulk of the restoration may be automated with the user specifying restoration-related parameters and a current region of the object image to restore.

5. Visual Compensation

To generate our final visual compensation, we use our restored image as the target appearance and our program computes the imagery necessary to project from each of the projectors. Using the multiple overlapping projectors in our setup produces a light-efficient, visual compensation which best alters the appearance of an object. We define an optimal alteration as achieving a desired balance between supporting the maximum light intensity and capable color variation while still producing a smooth visual compensation. Further, in order to limit damage due to prolonged exposure to light, the system constrains the amount of light intensity to be incident per unit surface area of the object below a fixed constant E_{bound} .

The amount of light that the object receives is particularly important when working with light-sensitive objects. In general, the user only needs to define E_{bound} and our program will balance out the projectors' light contributions to the visual compensation in such a way to achieve a visually appealing result.

Generating a light-efficient compensation

The efficiency with which any one projector can change the color and intensity of an object varies with respect to the object's surface orientation relative to the projector and the object-projector distance. As the object's surface orientation reaches a grazing angle with a projector's orientation, that projector would have to shine a greater amount of light (and thus more energy) onto the object to achieve the equivalent brightness from a projector oriented head-on with the object's surface. Similarly, since light energy dissipates as it travels through air, a projector positioned closer to the object's surface will yield a brighter compensation with less light than a projector positioned further away.

These two factors suggest that a proper balance of light contribution from each of the system's projectors (and further, each of their individual pixels) leads to both a smooth and bright compensation while respecting E_{bound} . We achieve this goal by computing the weights for each pixel in each projector. Each pixel's weight considers the distance and orientation of the projector pixel to the surface point it is illuminating. If the pixel is well-positioned and well-oriented, then the pixel's weight will be high. If not, the pixel's weight will be low.

With the proper weights calculated for each projector pixel, the only task remaining is to generate a compensation image for each projector. This is done by taking the previously restored image of the object and, with the help of a radiometrically calibrated function which converts a color intensity value viewed by the camera to the input intensity for a projector pixel, computing the colors necessary to shine onto the object itself to achieve the desired appearance.

Virtual illumination

Our system optionally supports virtual illumination of the object with a set of virtual light sources. Virtual illumination can be used for a variety of purposes such as accentuating an important area of an object. In addition, a small amount of the noise from the original image can be infused back into the restored image so as to simulate some of the irregularities and material properties of the original object. This effect is done by applying a lighting scheme onto the 3-D object in our program and using this result as our target appearance for the projectors to display.

6. Results

We have applied our system to a group of objects borrowed from the Indianapolis Museum of Art's (IMA) study collection and Indiana University-Purdue University Indianapolis' Anthropology Department. In particular, the objects in Figures 3 and 6 are in the style of ceramic ware from the Casas Grandes region in northern Mexico, the object in Figure 4 is an accurate replica of a Chinese vessel from the Neolithic period (2000-2500 B.C.), and the object in Figure 5 is a replica of a ceramic vessel from the Moche Culture in northern Peru.

For all objects, the preprocessing work was to securely place the object on the platform, take a photograph under diffused illumination, perform a self-calibrating 3D acquisition, compute the light-transport matrix, and carry out a radiometric calibration. Preprocessing is automated and takes about 1-2 hours to complete per object, with half of the time spent acquiring images.

Using our computer program, the user can restore the object to a plausible original appearance and generate views for different virtual illumination setups. For the restoration, the user specifies a few intuitive numeric parameters to specify the amount of smoothness and restoration desired. For the computation of the compensation image, the user selects the target luminance energy E_{bound} . The restoration results in this paper were created in 5-30 minutes.



Fig. 3: Steps taken during the virtual restoration of a ceramic vessel in the Casas Grandes style. a) Photograph of the diffusely-illuminated original object used for color classification. b) Image of the synthetic restoration produced by our system, which has been infused by a small amount of noise and mild virtual illumination. c) Photograph of the virtually restored object based on b) and under projector illumination.



Fig. 4: Steps taken during the virtual restoration process for a replica Neolithic Chinese vessel. a) Photograph of the diffusely-illuminated original object. b) Image of the synthetic restoration. c) Photograph of the virtually restored object based on b). The insets show a before and after close-up of the top portion of the vase which demonstrates our system's ability to complete the cross-hatch pattern as part of the restoration process.

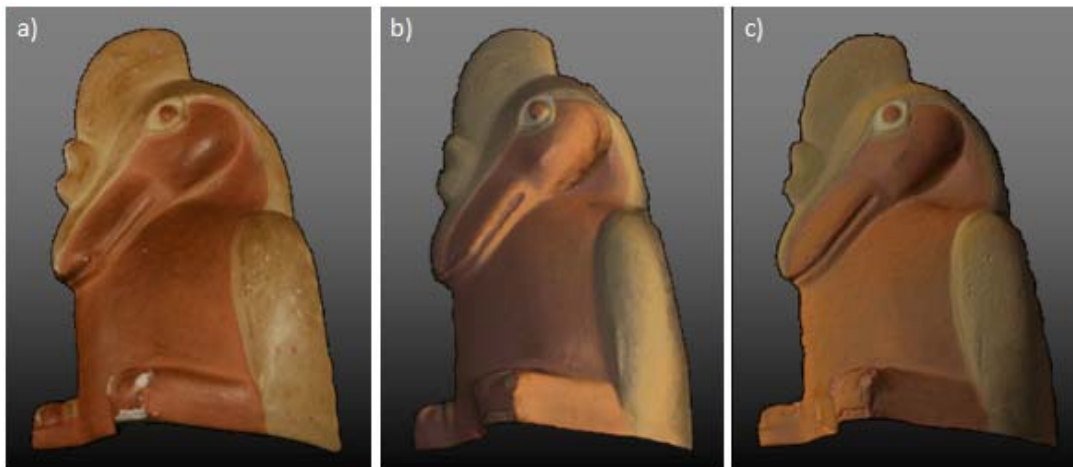


Fig. 5: Virtual illumination of a replica of a ceramic vessel from the Moche Culture in northern Peru. a) Photograph of the original object. b) Photograph of the virtually restored object with virtual illumination (high specularity). c) Photograph of the virtually restored object similar to b) but with a more diffuse virtual illumination. Note, the handle of the vessel could not be well captured so it was removed from the image.




Fig. 6: Virtual restoration of a ceramic vessel in the Casas Grandes style. a) Photograph of the diffusely-illuminated original object. b) Photograph of the virtually restored object. c) Detail image of surface irregularities and losses. d) Close-up of the same area on the virtually restored object.

Limitations

Our system does have some limitations. During image restoration, the gamut of color patterns for our color classification scheme must be relatively small and finite. We do not handle gradients of colors, as these patterns contain such a large number of colors that pixel selection and thus color classification becomes infeasible. Furthermore, the ability to control the user-defined parameters effectively is limited when dealing with complex patterns (such as human or animal figures) which in turn makes restoration difficult. Mild specularity on the object's surface is supported by our acquisition and compensation method but in general is difficult to handle. Since it is difficult to visually alter objects with subsurface scattering (for example, those comprised of fabric, feathers, leather, and glass) our current system is unable to virtually restore them.

7. Conclusions and Future Work

Our virtual restoration system enables us to alter the appearance of objects to that of a synthetic restoration, to create virtual re-illuminations of them, and to achieve a balance of high brightness and contrast with smooth visual compensation. Our approach provides a physical inspection stage where viewers are not restricted to viewing a virtual restoration on a computer screen, but rather can see and experience the restored object in person.

Now that the model has been created, we are seeking feedback from a broader audience on the utility of this type of system in a museum setting. Two possible scenarios that we have explored are:

1. Using this system in a conservation lab to enable conservators, curators, and other stakeholders to envision potential results of actual conservation treatments.
2. Using this system as an interactive display device in museum settings. The system could be presented as a stand-alone unit with just one object or incorporated within an exhibition. By providing multiple virtual restorations, the system provides the viewer the opportunity to choose different presentation states of the same object.

We would like to extend our method to support multiple-camera setups and smaller projectors. This would enable us to capture a more complete geometric model, to perform a visual compensation over a wider field of view, and to have a more compact hardware system. For instance, because our current field-of-view is limited, the handle of the vessel in Figure 5 could not be well captured and thus we removed it from the visual compensation image. Smaller and cheaper projectors can be used and the additional artifacts (e.g., stronger vignetting) can be handled by the radiometric calibration. We are also interested in developing a gamut of interactive restoration operations to be performed on the observed objects using a camera-based feedback loop. Finally, we are pursuing portable virtual restoration stations.

Acknowledgements

This research is supported by NSF CCF 0434398 and by an Indiana/Purdue University Intercampus Applied Research Program Grant.

References

- Aliaga, D., Law, A.J., and Yeung, Y.H. (2008). A Virtual Restoration Stage for Real-World Objects. *Proc. ACM SIGGRAPH Asia*, ACM Transactions on Graphics, 27, 5.
- Aliaga, D., and Xu, Y. (2008). Photogeometric Structured Light: A Self-Calibrating and Multi-Viewpoint Framework for Accurate 3D Modeling. *Proc. of IEEE Computer Vision and Pattern Recognition*, 1-8.
- Koller, D., Trimble, J., Najbjerg, T., Gelfand, N., and Levoy, M. (2006). Fragments of the City: Stanford's Digital Forma Urbis Romae. *Proc. Of the Third Williams Symposium on Classical Architecture, Journal of Roman Archaeology Suppl. 61*, 237-252.
- 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. (2000). The Digital Michelangelo Project: 3D Scanning of Large Statues. *Proc. ACM SIGGRAPH*, 131-144.
- Nautiyal, V., Clark, J.T., Landrum III, J.E., Bergstrom, A., Nautiyal, S., Naithani, M., Frovarp, R., Hawley, J., Eichele, D., Sanjiv, J. (2003). 3D Modeling and Database of Indian Archaeological Pottery Assemblages - HNBSU-NDSU Collaborative Initiative. *Proc. Computer Applications and Quantitative Methods in Archaeology Conference*, 8-12 April 2003, Vienna, Austria.
- Raskar, R., Welch, G., Low, K.L., and Bandyopadhyay, D. (2001). Shader Lamps: Animating Real Objects With Image-Based Illumination. *Proc. Eurographics Workshop on Rendering Techniques*, 89-102.
- Skousen, H. (2004). Between Soil and Computer – Digital Registration on a Danish Excavation Project. *British Archaeological Reports International Series*, v. 1227, 110-112.
- Viti, S. (2003). Between Reconstruction and Reproduction: The Role of Virtual Models in Archaeological Research. *Proc. Computer Applications and Quantitative Methods in Archaeology Conference*, 8-12 April 2003, Vienna, Austria.