

A Virtual Restoration Stage for Real-World Objects

Daniel G. Aliaga

Alvin J. Law

Yu Hong Yeung

Purdue University



Figure 1. Virtual Restoration. We present a system to restore and alter the physical appearance of old and deteriorated objects: a) photograph of an ancient Chinese vase (ca. 2250 BC); b) an image of the synthetic restoration produced by our algorithm; c) photograph of the restored object using our system and guaranteed not to exceed a provided maximum amount of light per unit surface area (note: in photographs, background is illuminated by a light source behind the object causing shadows on the rim of the object).

Abstract

In this paper, we introduce a system to virtually restore damaged or historically significant objects without needing to physically change the object in any way. Our work addresses both creating a restored synthetic version of the object as viewed from a camera and projecting the necessary light, using digital projectors, to give the illusion of the object being restored. The restoration algorithm uses an energy minimization method to enforce a set of criteria over the surface of the object and provides an interactive tool to the user which can compute a restoration in a few minutes. The visual compensation method develops a formulation that is particularly concerned with obtaining bright compensations under a specified maximum amount of light. The bound on the amount of light is of crucial importance when viewing and restoring old and potentially fragile objects. Finally, we demonstrate our system by restoring several deteriorated and old objects enabling the observer to view the original or restored object at will.

CR Categories: I.3 [Computer Graphics], I.3.3 [Picture/Image Generation], I.3.5 [Computational Geometry and Object Modeling], I.3.6 [Methodology and Techniques].

Keywords: digitization, restoration, image completion, energy minimization, radiometric calibration, light transport.

1 Introduction

The objective of our work is to visually restore deteriorated and historically-significant objects viewed only by the naked-eye

without needing to touch or change the original object in any way. Current physical restoration and conservation treatments that are normally performed are very subtle because the owners of the objects do not want them to be significantly changed. Conservation treatments are usually limited to the minimum necessary to prevent further deterioration. Restoration efforts are faced with the dichotomy of keeping the object original and of providing a view of the object as it was initially intended to be seen. Further, a historical artifact might actually have multiple possible restorations, either because of uncertainty in its former appearance or because the artifact appeared differently depending on the time period. Our work circumvents these issues and limitations by altering the *visual appearance* of the objects in a flexible and controlled manner. This allows museums, for instance, to resolve the dichotomy and to showcase a genuine historical artifact with different appearances.

Our methodology computes a restored form of the observed object and uses several digital projectors and a digital camera to alter the appearance of the object by precisely controlling the position, color, and intensity of projected light. Nevertheless, several other approaches to restoration exist. A user could simply use an image editing program (e.g., Photoshop) to restore a photograph of the object. However, the use of such programs may require many hours of arduous detailed work whereas our restoration process can be completed in minutes. The restored view of the object can be displayed on a computer screen providing significant flexibility but lacks the benefit of the natural cues of depth perception, parallax, and physical inspection. Physically-restored replicas can be fabricated but the task is time consuming and is limited to a single interpretation of the restored appearance of the object. Creating simple replicas (e.g., diffuse white replicas capturing the general shape) and using projectors to augment the replica with visual details does give the general object appearance, but unless all observers are tracked and provided with a custom image (e.g., special goggles, head-mounted display), the intricate and view-dependent details of a physical object cannot be provided. Moreover, museums pride themselves on showcasing genuine

artifacts. With our methodology, a restored view is interactively computed within minutes and many simultaneous viewers can look at the actual historical object in a virtually restored state, have the benefit of depth perception, parallax, and physical inspection, observe all the intricate details of the object, and not need any special goggles or viewing devices.

The main challenges lie in (i) correcting the deterioration of the colors on the object's surface, (ii) providing a restored visual appearance using no more than a user-specified maximum amount of light, and (iii) easily obtaining the needed information about the object's and projectors' geometric and radiometric properties. The targeted object might suffer from missing chips of paint, stains, color fading, and cracks that noticeably alter its current appearance and hinder inferring the object's original appearance. Additionally, a common concern for historically-significant objects that are on display is the amount of incident light. While a restoration scheme like ours could be enabled only periodically, we seek to maximize the visual restoration that can be obtained under a specified maximum amount of projected light. This allows the user to control the amount of light desired to be projected onto the object. Moreover, since the geometric and radiometric properties of the object and projectors are initially unknown, the restoration stage should be able to accurately and easily obtain this information so as to make deploying our restoration stage robust and practical.

Our two key observations are that using multiple projectors with fully overlapping fields-of-view and assuming a discrete number of colors for the targeted object enables a robust virtual restoration system using a user-specified maximum amount of projected light. The first observation provides us with the flexibility to choose at the pixel-level how to best combine the projectors in order to perform the visual compensation within a specified maximum amount of light. The second observation hinges on the fact that our targeted historical and aged objects are often hand-painted and consist of a small number of foreground and background colors. The colored patterns on the surface have deteriorated which causes the observed color discontinuities to no longer be the contours of the painted shapes. Unlike image filtering (e.g., edge-preserving bilateral filtering), this observation enables us to explicitly change the contours of the deteriorated patterns. Further, as opposed to image inpainting efforts, our approach is not limited to filling in small holes, does not depend on self-similarity within the image, and is able to tackle images with significant structure. By assuming few colors are on the object's surface, we devise a robust and accurate color classification method. The classification converts the problem into a discrete optimization framework and enables us to iteratively alter the shape of the color patterns to attempt to satisfy a set of smoothness criteria. Furthermore, the restoration optimization only allows a lighter luminance pixel to be replaced with a similar or darker luminance pixel. This restriction significantly improves the visual quality of restoring a target object with projected light.

Overall, our approach uses an object stage and three main steps to perform a virtual restoration. The stage consists of a platform for the object as well as three digital projectors and a digital camera pointing towards the object. During an initial acquisition step, a self-calibrating 3D reconstruction method obtains a dense model of the observed object along with the poses of the projectors and camera. The high accuracy and self-calibrating nature of this step are necessary to ensure ease of use and to provide visually-compelling restorations. Second, an interactive *image restoration method* uses a photograph of the object to produce a restoration image which attempts to infer an original appearance of the

object. This step assigns each pixel to one of a small number of dynamically chosen colors using our color classification algorithm. Then, each patch of same-color pixels is restored using an energy minimization framework which performs a sequence of contour smoothing operations. Each operation attempts to reduce the maximum curvature of the contour while respecting inferred spatial relationships with nearby contours and producing a compensation-complaint restoration of the color patterns. Third, a *visual compensation algorithm* is used to compute compensation images for simultaneous projection. Our algorithm uses the results from the initial acquisition step to obtain a precise estimate of the light transport between projectors and cameras (e.g., similar to [Sen et al. 2005]). Then, our method uses a surface radiance model to calculate projector images so as to best alter the visual appearance of the target object to that of the restored image while guaranteeing not to exceed a given maximum amount of light per unit surface area for the object.

Altogether, we present a complete and self-calibrating system requiring only a small amount of user input to perform compelling virtual restorations of real-world objects. Our system uses only off-the-shelf hardware. The result is a fundamentally different way of viewing and understanding important and perhaps historically-significant objects. We demonstrate our system by performing virtual restorations on a variety of historical objects exhibiting varying degrees of deterioration.

Our major contributions are

- a first-of-its-kind system that enables any number of simultaneous viewers to view in place a deteriorated real-world object either in its original state or in a restored state,
- a restoration method to infer an original appearance of a target object and to produce an image of its restoration, and
- a visual compensation algorithm that uses a surface radiance model and light-transport based approach to alter the target object to its restored version using no more than a user-specified maximum amount of light per unit surface area.

2 Related Work

Our research builds upon work in altering the visual appearance of objects using digital projectors, computing an efficient radiometric compensation, and performing an image-based restoration of the observed object. Raskar et al. [2001] present 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 a simple photometric calibration. Subsequently, several radiometric compensation algorithms have been proposed for calibrating projectors and cameras [Mitsunaga and Nayar 1999; Grossberg et al. 2004; Fujii et al. 2005; Wetzstein and Bimber 2007] often for the purpose of projecting movies on top of arbitrary surfaces. Grossberg et al. [2004] did alter the appearance of an object but assumed a very simple diffuse object model. Wetzstein and Bimber [2007] used a light-transport approach similar to ours for capturing arbitrary surface properties. Augmented reality frameworks provide both real and synthetic content (e.g., [Azuma et al. 2001; Bimber et al. 2001]) but often need head-mounted displays or support a limited number of viewers. None of these methods calculate how much light is used during compensation or address visual restoration. Our work significantly extends these general concepts to enable observing and restoring fragile and deteriorated objects.

With regards to visually repairing the observed object, image restoration methods can be loosely divided into algorithms for image denoising and for image inpainting. Many image denoising and filtering methods have been proposed to restore the quality of captured images and, in general, address removing very small artifacts. Since the sharp discontinuity of edges is often desired (e.g., the edges of a solid-colored pattern painted on the object), edge-preserving smoothing methods are particularly relevant to our objective; e.g. anisotropic diffusion [Perona and Malik 1990] and bilateral filtering [Tomasi and Manduchi 1998; Durand and Dorsey 2002; Weiss 2006]. However, in our case, the observed edges are not necessarily the ones we want to keep. Our objective involves more significant structural changes to the observed patterns since missing chips of paint, blemishes, and cracks will cause undesired contours to appear. Thus, our task must include adjusting the contours to a more correct shape and location.

Image inpainting is concerned with “filling-in” a user-selected region. A first strategy is to grow the region surrounding small holes [Bertalmio et al. 2000; Chan and Shen 2001; Levin 2003]. A second approach is to incrementally fill in larger holes by searching for patches of pixels elsewhere in the image (or video in some cases); e.g., [Drori et al. 2003; Criminisi 2003; Jia and Tang 2003]. A third approach is to merge with texture synthesis approaches (e.g., [Efros and Leung 1999; Wei and Levoy 2000]) and to generate approximate content for the missing regions (e.g., [Igehy and Pereira 1997; Bornard et al. 2002; Bertalmio et al. 2003]). While these methods produce very compelling results and are able to complete regions with large omissions, they have difficulty with structured images, assume the region to inpaint is surrounded by valid image content, and need an a priori mask specifying the location of the holes to inpaint. In our case, the objects typically contain hand designed patterns with deterioration throughout the entire surface. Using inpainting to reproduce the structure is challenging since no holes are explicitly present. Sun et al. [2005] present an image completion method which improves upon structure-oblivious inpainting but requires user-specified structural hints for all structural corrections.

In contrast, our restoration approach is more related to the object-editing methodology of Barrett and Cheney [2002] which enables altering an image at the object-level rather than at the pixel-level. In a similar spirit to Barrett and Cheney [2002] (and also to Sun et al. [2005]), our method asks the user to provide hints for steering the restoration towards an original structure of the object. However, our algorithm uses a color classification algorithm to convert the problem into a discrete optimization. This enables automatically handling the details of inflating or deflating the painted patterns, filling in small holes, and ensuring a set of smooth contours while maintaining the spatial relationship between them. While general color classification methods can partition pixels into a set of similar-colored clusters (e.g., [Comaniciu et al. 2002; Felzenszwalb and Huttenlocher 2004]), we use a dynamically calibrated k-means clustering approach specialized towards the case of knowing only a few colors are present. This method converts the problem to a discrete optimization and significantly reduces the effort on part of the user to specify regions in need of restoration.

3 Object and Acquisition Stage

Our virtual restoration stage is intended to have an object on display so that either its original or restored state can be viewed at the press of a button. For example, a computer-controlled timer can enable the restoration system for no more than a total exposure time. Nevertheless, our visual compensation system ensures a maximum amount of light (e.g., expressed as a



Figure 2. Object Stage. The stage used by our approach consists only of self-calibrated off-the-shelf hardware.

percentage of the projector’s maximum illumination strength) is not exceeded and thus prevent damage to the object over time.

The restoration session for the object starts by (i) placing the object on the stage, (ii) acquiring a geometric model of the object and estimating the internal and external parameters of the projectors, (iii) capturing a picture for computing the image-based restoration, and (iv) efficiently calculating the light-transport matrix and radiometric calibration. Then, a compensation image for each projector is interactively computed using a desired set of restoration parameters and maximum light-intensity. Figure 2 contains a picture of the setup with an example object in place.

To facilitate easy deployment of our stage, we perform a self-calibrating reconstruction of the object and of the projector parameters. Our 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 poses and focal length of the projectors. The result is a dense and accurate model containing position and normal data of the target object.

4 Restoration

Our image restoration method provides an interactive application whereby a user can restore the appearance of the observed object using an iterative energy minimization process. Figures 3a and 3b show an example object section highlighting the deterioration. Our goal is to arrive at Figure 3c, a plausible restoration of the highlighted object section, and Figure 3d, which shows Figure 3a restored. Our method is able to work at any scale so as to provide both minor and major alterations to object appearance. The user interactively controls the smoothness of the restoration as well as how much it can deviate from the initial appearance.

To begin the restoration process, our color classification algorithm first assigns each pixel to one of a small set of colors. Then, the user selects a *region* of the image for restoration and assigns the region a background color. The size of the region depends on the




Figure 3. Image Restoration. a) A close-up image of a deteriorated object. 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).

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. These criteria, 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.

4.1 Color Classification

Using an image of the object as the input, our classification places each pixel into one of a small set of colors by using a calibrated k-means clustering method. The image to be used for restoration is captured using indirect and diffused illumination so as to reduce the troublesome 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.

The calibration pixels are used in an optimization process for defining a color distance function that provides maximum separation between the $\binom{n}{2}$ pairs of representative colors on the object. Calibration pixels corresponding to the same color are averaged to produce a set of n representative colors of the object. Generically, the distance function between two colors c_i and c_j , expressed in the xyY color space, is

$$D(c_i, c_j) = a_1(c_{ix} - c_{jx})^2 + a_2(c_{iy} - c_{jy})^2 + a_3(c_{iY}, c_{jY})^2 \quad (1)$$

where a_1 , a_2 , and a_3 are positive weights given to each channel and $a_1 + a_2 + a_3 = 1$. These weights balance the effect of chroma and intensity differences between the representative colors. Using a nonlinear least squares optimization (e.g., Levenberg-Marquardt), we compute values for a_1 and a_2 by minimizing the inverse of the distance between all pairs of representative colors and define $a_3 = 1 - a_1 - a_2$. The distance function is then used to classify each pixel on the image to one of the n representative colors using k-means clustering.

4.2 Energy Minimization

Our energy minimization 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.

Figures 4a-b motivate these restoration criteria using an example object area. 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 physical object, ensuring the restoration can be done well via radiometric compensation is crucial (Figures 4c-d).

Altogether, the previous criteria are used in an optimization to iteratively reduce the maximum pixel energy term for all pixels in all contours to below a threshold $\alpha > 0$. The energy term

$$e_{ij} = \kappa(p_{ij})R(p_{ij}, \hat{p}_{ij})B(p_{ij}, \hat{p}_{ij}) \quad (2)$$

corresponds to the i 'th pixel on the closed poly-line curve defining contour j . The functions $\kappa(p_{ij})$, $R(p_{ij}, \hat{p}_{ij})$, and $B(p_{ij}, \hat{p}_{ij})$ represent the three restoration criteria, respectively. \hat{p}_{ij} is the proposed new position for p_{ij} defined by a Laplacian smoothing operator applied to the contour points, namely

$$\hat{p}_{ij} = \frac{p_{i-1j} + p_{i+1j}}{2} \quad (3)$$

All contour points p_{ij} are placed in a priority queue, and for each iteration the point with the largest e_{ij} value is selected. Once the point p_{ij} has been moved, the algorithm recomputes the normal of the contour segment before and after p_{ij} as well as the now changed energy term e_{ij} . Neighboring contour points must also have their normals and e_{ij} values recomputed since these values are a weighted combination of the point's and its neighbors' normals and e_{ij} values. Termination is reached once all $e_{ij} \leq \alpha$ where α controls the amount of smoothing. As α decreases, more smoothing occurs. We typically use α values ranging from 1 to 8, with 2 being a reasonable value for most cases.

While the contour points are initially defined to be at exact pixel locations, they can move to fractional pixel locations during the minimization. Thus, before minimization the contours are




Figure 4. Restoration Criteria for the Energy Minimization. a-b) Contour smoothing and contour repelling/attracting restores the overall spatial relationship of the contours on the whole. c-d) Regions circled in green demonstrate compensation-compliant contour changes while regions circled in red do not. In all, the dotted lines represent the original shape of the contours.

removed from the image and replaced with the background color. Then, after minimization the changed contours are rasterized back onto the image and over the background color. For efficiency, if the distance between the contours points before and after p_{ij} is less than a small threshold (e.g., we typically use a distance of 0.2 pixels), then p_{ij} is removed altogether from the contour.

4.2.1 Contour Smoothness

Obtaining a smooth contour is done by reducing the maximum curvature $\kappa(p_{ij})$. The curvature κ at a contour point p_{ij} is defined as the change in angle between the previous and next contour segments' normals over the arc length of the two segments:

$$\kappa(p_{ij}) = \frac{\frac{v_{ij}}{\|v_{ij}\|} \cdot \frac{v_{i+1j}}{\|v_{i+1j}\|}}{(\|v_{ij}\| + \|v_{i+1j}\|)} \quad (4)$$

where $v_{ij} = p_{ij} - p_{i-1j}$. Since the contour points are initially created from discrete pixels, only a small number of initial curvature values exist. As a consequence, there is little distinction between pixels near and far from an area of high curvature. To obtain a more continually varying measure of curvature, we smooth the normals of the segments with those of their neighbors. The user controls the magnitude of contour smoothing by specifying the number of neighbors to use. The larger the number of neighbors the user selects, the greater the degree of smoothing.

4.2.2 Contour Repelling/Attracting

Our method restores the overall spatial relationship of the contours by attempting to move each contour pixel to a computed distance from its corresponding closest pixel on an opposing contour. This procedure effectively either attracts or repels contour pixels from each other until converging to a mutually agreed upon configuration (Figure 4a).

For point p_{ij} on contour j , we find the closest point q_{uv} on a neighboring contour v that is directly visible by p_{ij} . This pairing is not bijective because when q_{uv} finds its nearest neighboring point, it may not necessarily be p_{ij} . The distance $d_{ij} = \|p_{ij} - q_{uv}\|$ is defined as the *repelling distance* for a contour point p_{ij} . The repelling distances of adjacent contour points are combined via weighted averaging yielding a smoothly varying distance d_{ij}^* per contour point. It should be noted that a sufficiently close pairing point q_{uv} does not necessarily exist for all points p_{ij} . A direct line of sight between paired contour points is present if a ray cast from p_{ij} reaches q_{uv} uninterrupted by a contour segment. Since these computations are applied on the original contours in a discretized pixel space, visibility can be confirmed if the raster line between the two points does not cross a contour pixel which is neither p_{ij} nor q_{uv} .

The function $R(p_{ij}, \hat{p}_{ij})$ measures whether a change from p_{ij} to \hat{p}_{ij} brings a contour into closer agreement with the desired distance d_{ij}^* to an opposing contour. If we define \hat{d}_{ij} to be the distance between \hat{p}_{ij} and q_{uv} , then our goal is to perform $p_{ij} \rightarrow \hat{p}_{ij}$ if and only if $|\hat{d}_{ij} - d_{ij}^*| \leq |d_{ij} - d_{ij}^*|$. A function that captures all these objectives is

$$R(p_{ij}, \hat{p}_{ij}) = \begin{cases} \frac{1}{\beta(1 + |d_{ij} - d_{ij}^*|)} & \text{if } |d_{ij} - d_{ij}^*| < |\hat{d}_{ij} - d_{ij}^*| \\ 1 & \text{if } q_{uv} \text{ is undefined} \\ \beta(1 + |d_{ij} - d_{ij}^*|) & \text{if } |d_{ij} - d_{ij}^*| \geq |\hat{d}_{ij} - d_{ij}^*| \end{cases} \quad (5)$$

where $\beta \geq 1$ is an emphasis factor.

The variable β weighs the relative importance of a contour movement towards reaching a distance d_{ij}^* as compared to the rest of the terms in e_{ij} . If the user wishes to emphasize convergence towards d_{ij}^* , β should be assigned a high value, resulting in larger e_{ij} values when a smoothing operation moves closer towards d_{ij}^* . If not, low values for β will accentuate smoother contours while not necessarily pushing the contours to a distance value of d_{ij}^* . In general, β controls the amount of contour movement allowed, and for our objects a typical β value ranges from 1.5 to 8.

4.2.3 Compensation-Compliant Contour Changes

Our algorithm creates a compensation-compliant restoration image by encouraging a solution where pixels usually lower their intensity (i.e., luminance) when they are being modified (Figure 4b). In general, the compensation images will attempt to change the color of points on the object surface to that of the restoration image. The compensation process is calibrated and optimized to work on the whole. Nevertheless, it remains difficult to change the color of an object point whose spectrum and intensity of reflected light does not coincide sufficiently with the spectrum and intensity of the corresponding color in the restoration image. In the extreme case of an object point that only reflects light within a very narrow band, little color changing is possible. Fortunately, very few materials that reflect nearly a single-frequency are used by our typical target objects. Thus, the color of a point on the object can be changed to another one by changing the relative strength of the RGB components of each pixel of projector light. However, the most limiting factor is the amount of light that can be reflected back by a given point on the object.

To ensure contour changes are compensation-compliant, we allow a pixel to change its relative RGB values but only allow it to maintain or reduce its intensity. For this objective, we define

$$B(p_{ij}, \hat{p}_{ij}) = \begin{cases} 0 & \text{if } I(p_{ij}) > I(\hat{p}_{ij}) + \gamma \\ 1 & \text{if } I(p_{ij}) \leq I(\hat{p}_{ij}) + \gamma \end{cases} \quad (6)$$

where $I(p)$ is the intensity value of point p_{ij} and γ represents a small tolerance range allowing pixels to become slightly more intense. This guidance does affect the appearance of the restored image (as seen in Figure 4b). Nevertheless it is done simultaneously with the other metrics and thus a restoration can still be found that best maintains all criteria.

4.3 Acceleration

Since the energy minimization moves contours pixel by pixel, it is highly parallelizable as long as neighboring pixels are not processed simultaneously. Thus, to improve interactive processing time, we partition the current region of the image being restored into a set of $g \times h$ grid cells. The size of the cells should be small enough to process quickly yet large enough to capture feature details and their surrounding neighbors. Given a set of grid cells, we define a processing order to apply the energy minimization to multiple grid cells in parallel.

Our processing scheme is able to maintain both C^0 and C^1 continuity between patches and contours which span multiple grid cells. To ensure C^0 continuity between two grid cells, we anchor contour points which are touching a grid cell's perimeter. During the energy minimization process, we set the e_{ij} value of anchor points to be 0 so that they are always underneath α and therefore never moved. To ensure C^1 continuity, we expand the right and bottom sides of each grid cell by a percentage of the cell's width and height (e.g., we typically use 25%). This expansion allows us to modify parts of the neighboring grid cells to ensure a smooth transition from one cell to another. However, this expansion also

results in grid cells which overlap and can lead to race conditions when the same pixels are modified and indexed by two or more cells being processed simultaneously. To prevent race conditions, we process the cells in four non-overlapping groups, only processing the next group after the previous group is finished.

Each 2x2 set of grid cells are labeled following a standard left-to-right and top-to-bottom rasterization order. This yields the four groups *A*, *B*, *C*, and *D*, each with $gh/4$ cells, to be processed sequentially. The cells within each group are processed in parallel since they do not overlap. Altogether, an entire high-resolution image divided into grid cells can be processed within a few minutes or less depending on the amount of restoration needed.

4.4 User Hints

During image restoration, the user interactively provides hints to the system to guide large structural changes in the observed patterns due mostly to large missing chips of paint or severe color fading. Features which were originally part of the object may be too worn down to detect. In these cases, the user simply has to provide a rough sketch of the missing features on the image using a 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.

5 Visual Compensation

Our light-efficient visual compensation uses a surface radiance model and light from multiple overlapping projectors in order to best alter the appearance of an object. We define a best alteration as achieving a desired tradeoff between supporting the maximum amount of light intensity and color variation possible while still producing a smooth visual compensation. Further, in order to limit damage due to prolonged exposure to light, the tradeoff is subject to a constraint of a bounded amount of light intensity to be incident per unit surface area of the object. To perform this visual compensation, we use a formulation for light-efficient compensation, employ the light-transport from projectors to camera (and vice versa), and perform a radiometric calibration.

5.1 Light-Efficient Compensation

The efficiency with which any one projector can change the color and intensity of an object varies with three parameters: the object's surface orientation relative to the projector, the object-projector distance, and projector pixel resolution. In our formulation, we use the camera to obtain a sampling of unit surface area object points (i.e. they all have unit area). We use this sampling to ensure a bounded amount of light incident on the object per unit surface area. Since camera resolution is typically higher than projector resolution and both devices are at about the same distance from the object, this assumption is reasonable. Thus, for a projector *k* and a unit object point *i*, we define a set of smoothly changing compensation weights w_{ik} . Using these weights, we alter the appearance of the target object to that of the restoration image while not exceeding a desired maximum amount of light E_{bound} per unit surface area of the object. Further, we assume E_{bound} is less than or equal to the maximum light intensity per unit surface area output of one projector.

This overall objective can be expressed as trying to find the weights $w_{ik} \geq 0$ that solve for every unit object point *i*,

$$E_{bound} \geq E_i = \sum_{k=1}^{N_p} w_{ik} e_{ik} L_{ik} \quad (7)$$




Figure 5. Surface Radiance Model. The amount of light incident on a unit surface area object point from all projectors is modeled based on a diffuse surface illumination model.

where e_{ik} is the actual energy of the light ray from one pixel of projector *k* incident on object point *i*, L_{ik} is a measure of the portion of e_{ik} that is incident on the unit object point *i*, and N_p is the number of projectors.

5.1.1 Surface Radiance Model

We define a surface radiance model to measure the amount of light incident on a unit surface area object point from all projectors. Consider several $P_w \times P_h$ resolution digital projectors where each is emitting up to P_L lumens. Although a lumen is a measure of perceived light energy, it is often used for projectors. The maximum light energy emitted per projector pixel is approximately $e_p = P_L / (P_w P_h)$. We calculate an expression for estimating the portion of e_p incident on the object per unit surface area. We conservatively assume no loss of light energy occurs in the transmission from projector to object.

To compute the light energy per unit surface area, we assume a diffuse surface illumination model (Figure 5). The area of one projector pixel on the object's surface is approximately

$$A_{ik} \approx \left(\frac{\rho d_{ik}}{\cos \theta_{ik}} \right)^2 \quad (8)$$

where ρ is the per-pixel solid angle of the projector, d_{ik} is the distance between unit object point *i* and projector *k*, and θ_{ik} is the angle between the surface normal *i* and the light direction vector from object point *i* to projector *k*'s center-of-projection. This approximation is valid under the assumption that ρ is relatively small as compared to the distance d_{ik} . The maximum surface radiance from all N_p projectors onto a single object point *i* is

$$E_{i_{max}} = \sum_{k=1}^{N_p} e_p L_{ik} \quad \text{where} \quad L_{ik} = \frac{1}{A_{ik}} \quad (9)$$

Furthermore, given e_p and an object, the maximum light energy per unit surface area on the object is $\max_i(E_{i_{max}})$. Therefore, a reasonable range of values for E_{bound} is $[0, \max_i(E_{i_{max}})]$.

5.1.2 Combining Projector Lights

There are several ways to compute w_{ik} for each object point *i* and projector *k*, and each method achieves a different tradeoff between light intensity and visual smoothness. One naive option is to set all $w_{ik} = 1/N_p$ and to globally scale all projector pixels so that the total light energy from all projectors incident on any object point *i* never exceeds E_{bound} . While the threshold will not be exceeded, the overall light intensity of the visual compensation is diminished. The reduction occurs because energy is wasted on illuminating surface fragments at a grazing angle to a projector which are not lit as effectively as head-on illuminated object points. Further, areas not illuminated by all projectors appear abruptly darker than those covered by more projectors.

The most light-efficient weighting scheme is to choose for each object point i the single projector which *maximally* illuminates it. This corresponds to setting

$$w_{ik} = \begin{cases} 1, & L_{ik} = \max_j L_{ij} \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

in equation (7) and using equation (9) to compute values for L_{ik} . However, this has the side effect of causing shape discontinuities between the areas illuminated by different projectors because the projector selected to illuminate a smooth fragment of the object can abruptly change. To smooth across the boundary, we could linearly scale among all projectors so that the total amount of light energy emitted on any object point is bounded by E_{bound} . This is equivalent to calculating the weights as

$$w_{ik} = \frac{L_{ik}}{\sum_{j=1}^{N_p} L_{ij}}. \quad (11)$$

But this option may use a significant amount of light energy to illuminate surface fragments at a grazing angle to projector light. Because of the area-based scaling, the compensation image will be brighter than with uniform intensity per projector but not as bright as with the maximal efficiency configuration. Nevertheless, this linear scheme does provide a smooth transition between the intensities projected onto object points from different projectors.

Our solution achieves a parameterized tradeoff between the smoothness and energy efficiency by generalizing the weights as

$$w_{ik} = \frac{(L_{ik})^m}{\sum_{j=1}^{N_p} (L_{ij})^m} \quad (12)$$

where m can be any nonnegative number. On one hand, a larger m results in higher energy efficiency. On the other hand, a smaller m increases the visual smoothness of the compensation. The aforementioned schemes (e.g., uniform ($m = 0$), linear ($m = 1$), and maximum ($m = \infty$) distributions) are three particular cases of the generalized equation (12).

5.1.3 Calculating the Compensation Image

Given a desired set of weights to combine the light from multiple projectors, we must compute the e_{ik} 's to generate a visual compensation to make the object appear as in the synthetic restoration image. Let G correspond to a single-channel image of the RGB restoration image. The pixels of G map one-to-one with the set of sampled object points; thus pixel G_i corresponds to the intensity value for object point i . The projector pixel intensity value to yield an apparent value of G_i on the object's surface is

$$e_{ik} = s \cdot F_k(G_i) \quad (13)$$

where $F_k(*)$ is the radiometrically calibrated function which converts a color intensity value in G to an output intensity for projector k 's pixel incident on object point i and s is a global scaling factor to ensure E_{bound} is not exceeded per unit area on the object. Since several object points may fall incident on one projector pixel, we average all e_{ik} 's that map to the same projector pixel. In general, the functions F_k are nonlinear and are modeled via a radiometric calibration step (Section 5.2). To compute the global scaling factor, we first evaluate equation (13) with $s = 1$ to obtain the values E_i and then re-compute e_{ik} using

$$s = \frac{E_{bound}}{\max_i(E_i)}. \quad (14)$$

The final expression enabling us to perform a light-efficient and smooth compensation using a restoration image G and a light energy bound of E_{bound} is

$$E_{bound} \geq E_i = s \sum_{k=1}^{N_p} \frac{(L_{ik})^m}{\sum_{j=1}^{N_p} (L_{ij})^m} F_k(G_i) L_{ik}. \quad (15)$$

where m controls the amount of smoothness. In our experiments we typically use values for m in the range of 5 to 10.

5.2 Light-Transport from Projectors to Camera

To model the light interaction between the projectors and the camera, we compute a light-transport matrix and perform a radiometric calibration. While we obtain correspondence information between the projectors and cameras during the 3D acquisition phase, we do not obtain a precise estimate of the sizes and shapes of the regions of influence each projector pixel has on the object. Although the light-transport matrix can capture arbitrary surface reflectance properties, we assume the surface to be Lambertian and to exhibit only local diffuse inter-reflections. This enables us to significantly accelerate the process of acquiring the light-transport matrix. Similar to Sen et al. [2005], we acquire the transport for different parts of the object in parallel by shifting a grid of projector "dots" over the surface of the object. The dots are shifted to one of 16×16 or 32×32 positions, for example. Thus, in 16^2 or 32^2 images we can sample the per-pixel light-transport matrix for each projector. In our system, we add an additional refinement step whereby the initial correspondence between projector and camera pixels (i.e., from the 3D acquisition phase) is further refined during light-transport acquisition. The projection of each projector dot onto the object is typically a bright elliptical blob of pixels well separated from the neighboring blobs. Hence, for each projector dot we perform an image-search near the initially estimated corresponded camera pixel in order to snap to the center of the observed elliptical blob.

While the light transport implicitly captures the surface reflectance properties, we explicitly perform a radiometric calibration directly on the object itself in order to model the nonlinear behavior of the projectors (i.e., functions F_k in Section 5.1.3). We use a version of the method of Nayar et al. [2003] to calibrate each projector. Since our goal is to compensate for imagery relatively similar to the existing object, we assume a diagonal color mixing matrix and capture 255 gray-level images of the object as illuminated individually by each projector.

5.3 Virtual Illumination

We optionally augment the compensation with virtual lights shining on the object and/or with some of the (material) noise of the original photograph added back. To obtain a virtually illuminated compensation image, we apply the restored image as a texture-map on top of the 3D reconstruction of the object and synthetically render the object from the calibrated point of view of the camera. The rendering process can use any combination of OpenGL lights, and thus by reading back the frame buffer, we obtain a virtually re-illuminated version of the restored image. 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.

6 Results

We show the results of using our system to restore several deteriorated objects from various locations. Additional visualizations and live video recordings of the objects are included in the video accompanying this paper. We use three 1400x1050 Optoma EP910 projectors to restore the objects and one Canon Digital Rebel XTi 10MP camera to capture all photographs. The objects we used consist of genuine historical artifacts and replicas borrowed from the Indianapolis Museum of

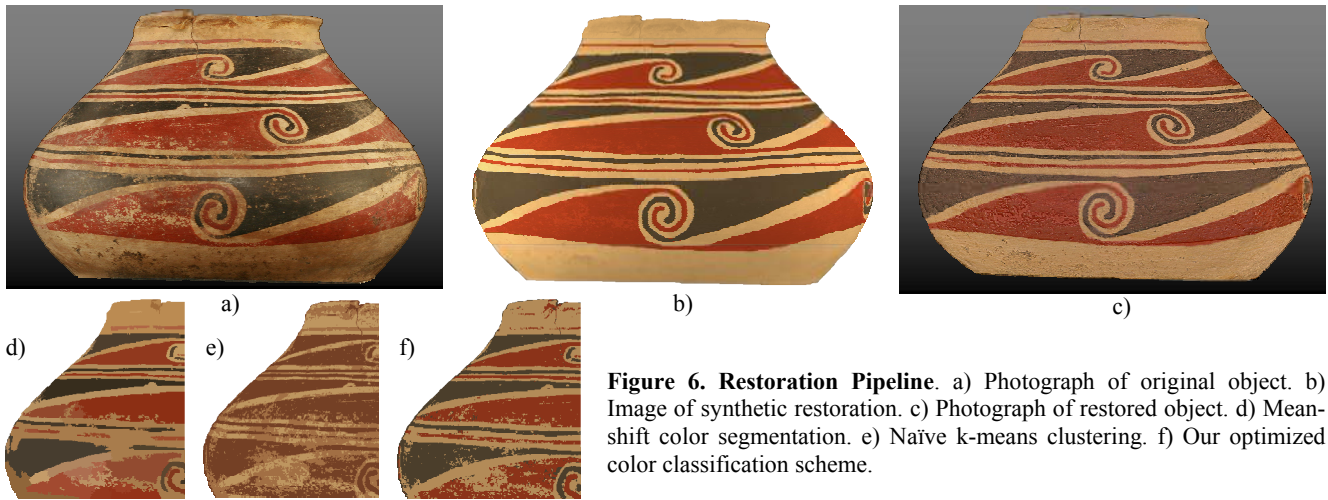


Figure 6. Restoration Pipeline. a) Photograph of original object. b) Image of synthetic restoration. c) Photograph of restored object. d) Mean-shift color segmentation. e) Naïve k-means clustering. f) Our optimized color classification scheme.

Art (IMA) and from the Eiteljorg Museum (both in Indianapolis, IN, USA). In particular, the objects in Figures 6 and 11 are genuine pottery from the Casas Grandes Cultural Region (1200-1425 A.D.) in northern Mexico, the object in Figure 1 is an accurate replica of a Chinese vase from the Neolithic period (2000-2500 B.C.), and the object in Figure 10 is a replica of a figurine from the Moche Culture in northern Peru.

For all objects, the preprocessing work was to securely place the object on the platform, take a picture 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 tool, the user can restore the object to a plausible original appearance and generate views for different virtual illumination setups. For the restoration, the user selects a value for the target level of smoothness α (as well as the number of adjacent pixels used to smooth the contour) and a value for β which limits the maximum amount of contour movement. For the computation of the compensation image, the user selects the target luminance energy E_{bound} and the smoothness exponent m . The restoration results in this paper were created in 5-30 minutes. Figure 6 shows example pictures produced during a restoration

session. Figure 6a contains a photograph of the diffusely-illuminated original object used for our color classification. Figure 6b shows an image of the synthetic restoration produced by our method. We have infused a small amount of noise and mild virtual illumination (as per Section 5.3). Figure 6c is a photograph of the restored object under projector illumination. All restorations are visible on the object by the naked eye and the observer is able to walk up to and around the object.

Our approach to color classification is tuned to images with few colors and outperforms more general methods. For instance, classification was performed using a publicly available implementation of the mean shift algorithm [Comaniciu and Meer 2002] (Figure 6d) and using a standard k-means clustering algorithm (Figure 6e). Neither of these approaches performed as well as our method (Figure 6f).

During interactive restoration, the user employs our program to generate a restored view of the object. For instance, Figure 7a shows a close-up of the initial classified pattern of the pictured object. Setting α to a small value causes an aggressive smoothing to occur during restoration (Figure 7b). Setting a large β value prevents contours from changing their current spatial relationship (Figure 7c). A good reconstruction is obtained by choosing a balance of these two parameters (Figure 7d). If the criterion to enforce a compensation-complaint restoration is removed, then an initial surface fragment such as the one in Figure 7e might be attempted to be restored as in Figure 7f which is difficult to perform. On the other hand, using our criteria can steer the restoration to a different solution that is easier to accomplish with visual compensation (Figure 7g).

Our visual compensation process achieves a trade-off between maximum brightness and visual smoothness subject to an energy




Figure 7. Energy Minimization Parameters. a) Original color classification. b) Highly smoothed: α and β small. c) Rigid contours: β large. d) Final balance of α and β . e) Photograph of original object. f) Non-compensation complaint contour change was to attempt to erase dark lines in yellow boxes. g) Compensation-complaint restoration.




Figure 8. Visual Compensation Brightness/Contrast. a) Maximally-bright projector combination but with artifacts. b) Linear combination of projectors (smooth but dim). c) Our nonlinear combination and solution (see equation 15).

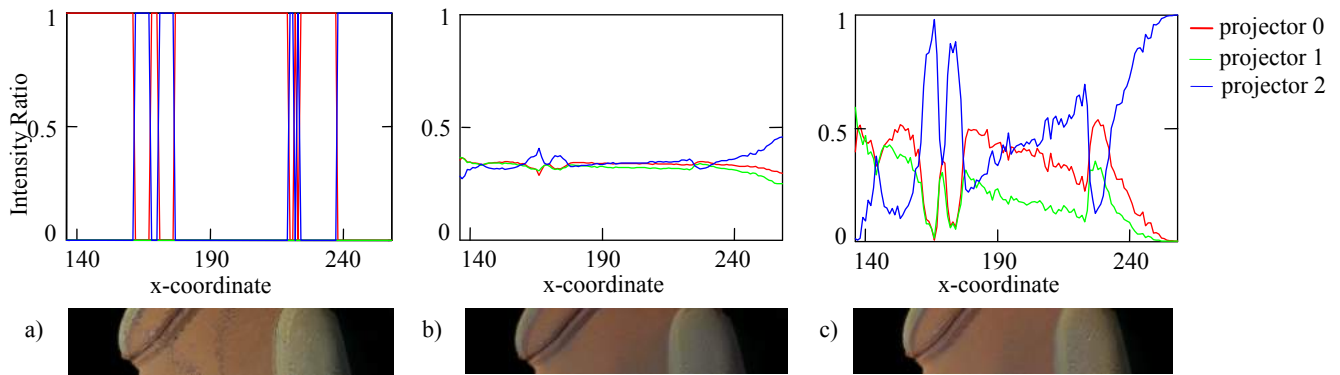


Figure 9. Projector Contributions. Graphs show the w_{ik} 's for a row of pixels across the middle of an image fragment. The sum under the area of for all projectors in each graph is the same: a) maximally-bright combination, b) linear combination, and c) our approach.

constraint. Figure 8a shows the maximally efficient compensation where each projector pixel chooses which object areas it best illuminates. The sudden changes between projectors may cause visible discontinuities on the object's surface. Using a linear combination achieves a smooth compensation image but at the cost of a reduction in the overall maximum brightness (Figure 8b). In contrast, by using a careful balance of lights, we achieve both high brightness and a smooth visual compensation.

Figure 9 numerically shows the weighting factors for the three projectors across a horizontal line of pixels from the middle portion of the restoration images shown. The weighting is such that the total area underneath the three curves sums to the same value for each of the three weighting schemes shown (Figure 9a: the projector pixels most head-on to an object point are enabled, Figure 9b: a linear combination of projectors, and Figure 9c: a non-linear weighted combination as per equation 15). The total area underneath each curve is limited by E_{bound} – essentially our method oscillates the curves so as to best fit to the bound.

Figures 10 and 11 show more examples. Figure 10 demonstrates two virtual illumination scenes. Figure 10a is a photograph of the original object. Figures 10b and 10c are photographs of the restored object under virtual lights. Figure 11 shows our system restoring an old pottery item. Figure 11a contains a photograph of the original object and Figure 11b is a photograph of the restored object under a very diffuse virtual light. The inset shows another photograph of the re-illuminated object using a different virtual light. The close-ups in Figures 11c and 11d show the visible condition of the object's surface before and after restoration. Our method is able to compensate for the deterioration and produce a restored appearance under bounded amount of light.

Limitations. Our techniques do have some limitations. During image restoration, the gamut of color patterns for our color classification scheme must be relatively small and finite. We do




Figure 10. Virtual Illumination. a) Photograph of original object. b) Photograph of restored object with virtual shiny appearance. c) Similar to (b) but a more diffuse appearance.

not handle gradients of colors, as these patterns contain such a large number of colors that pixel selection becomes infeasible. Furthermore, the ability to control the different parameters effectively is limited when dealing with complex items (such as human or animal figures) which in turn makes restoration difficult. Mild specularities are supported by our acquisition and compensation method but in general is difficult to handle. Our acceleration scheme for capturing the projector-camera light assumes at most local inter-reflection. Finally, we do not support objects with subsurface scattering since it is difficult to visually alter these types of surfaces with projector light.

7 Conclusions and Future Work

Our complete 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. The compensation is guaranteed to not exceed a specified maximum

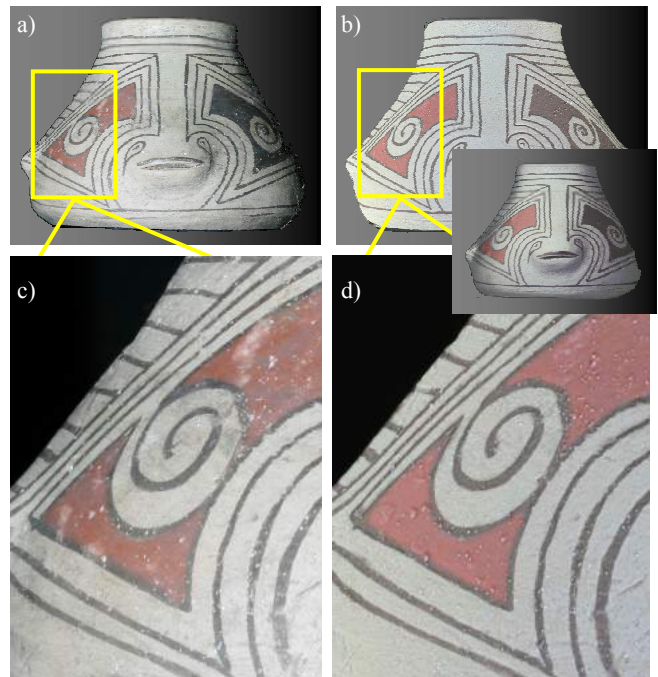


Figure 11. Additional Example. a) Photograph of original object. b) Photograph of virtually restored object (inset is another photograph of object under a different virtual illumination setup). c) Close-up of original deteriorations on object. d) Close-up of same area on virtually restored object.

amount of light per unit surface area. This requirement is critical for our approach to be deployed within museums and other locations which have a genuine interest in displaying original and restored versions of their artifacts. In such a setting, the amount of light exposed to the objects is of great concern. Our approach also 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”.

Future Work. First, 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 of our current limited field-of-view, the handle of the figurine in Figure 11 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. Second, we are 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.

8 Acknowledgements

We are grateful for the guidance provided by Larry Zimmerman, professor of Anthropology and Museum Studies at Indiana University, by Richard McCoy, Assistant Conservator at IMA, and by Amy McKune, Director of Collections at Eiteljorg Museum. We also thank the reviewers for their suggestions to improve this paper. This research is supported by NSF CCF 0434398 and by an Indiana/Purdue University Research Grant.

References

- ALIAGA, D., and XU, Y. 2008. Photogeometric Structured Light: A Self-Calibrating and Multi-Viewpoint Framework for Accurate 3D Modeling. In *Proc. of IEEE Computer Vision and Pattern Recognition*, 1-8.
- AZUMA, R., BAILLOT, Y., BEHRINGER, R., FEINER, S., JULIER, S., and MACINTYRE, B. 2001. Recent Advances in Augmented Reality, *IEEE Comp. Graphics & Applications*, 21, 6, 34-47.
- BARRETT, W. A., and CHENEY, A. S. 2002. Object-based Image Editing. *ACM Trans. on Graphics*, 21, 3, 777-784.
- BERTALMIO, M., SAPIRO, G., CASELLES, V., and BALLESTER, V. C. 2000. Image Inpainting. In *Proc. of ACM SIGGRAPH 2000*, 417-424.
- BERTALMIO, M., VESE, L., SAPIRO, G., and OSHER, S. 2003. Simultaneous Structure and Texture Image Inpainting. *IEEE Trans. on Image Processing*, 12, 8, 882-889.
- BIMBER, O., FROHLICH, B., SCHMALSTEIG, D., and ENCARNACÃO, L.M., 2001. The Virtual Showcase, *IEEE Computer Graphics & Applications*, 21, 6, 48-55.
- BORNARD, R., LECAN, E., LABORELLI, L., and CHENOT, J. H. 2002. Missing Data Correction in Still Images and Image Sequences. *International Multimedia Conference*, 355-361.
- CHAN, T. F., and SHEN, J., 2001. Nontexture Inpainting by Curvature-Driven Diffusions. *Journal of Visual Communication and Image Representation*, 12, 4, 436-449.
- COMANICIU, D., and MEER, P. 2002. Mean Shift: A Robust Approach Toward Feature Space Analysis. *IEEE Trans. Pattern Analysis Machine Intelligence*, 24, 5, 603-619.
- CRIMINISI, A., PEREZ, P., and TOYAMA, K. 2003. Object Removal by Exemplar-based Inpainting. In *Proc. of IEEE Computer Vision and Pattern Recognition*, 2, 721-728.
- DRORI, I., COHEN-OR, D., and YESHURUN, H. 2003. Fragment-based Image Completion. *ACM Trans. on Graphics*, 22, 3, 303-312.
- DURAND, F., and DORSEY, J. 2002. Fast Bilateral Filtering for the Display of High-Dynamic-Range Images. In *Proc. of ACM SIGGRAPH 2002*, 257-266.
- EFROS, A. A., and LEUNG T. K. 1999. Texture Synthesis by Non-Parametric Sampling. In *Proc. of IEEE International Conference on Computer Vision*, 2, 1033-1038.
- FELZENSZWALB, P. F., and HUTTENLOCHER, D. P. 2004. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59, 2, 167-181.
- FUJII, K., GROSSBERG, M. D., and NAYAR, S. K. 2005. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proc. of IEEE Computer Vision and Pattern Recognition*, 1, 814-821.
- GROSSBERG, M. D., PERI, H., NAYAR, S. K., and BELHUMEUR, P. N. 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *Proc. of IEEE Computer Vision and Pattern Recognition*, 1, 452-459.
- IGEHY, H., and PEREIRA, L. 1997. Image Replacement through Texture Synthesis. In *Proc. of International Conference on Image Processing*, 3, 186-189.
- JIA, J., and TANG, C. K. 2003. Image Repairing: Robust Image Synthesis by Adaptive ND Tensor Voting. In *Proc. of IEEE Computer Vision and Pattern Recognition*, 1, 643-650.
- LEVIN, A., ZOMET, A., and WEISS, Y. 2003. Learning How to Inpaint from Global Image Statistics. In *Proc. of IEEE International Conference on Computer Vision*, 1, 305-312.
- MITSUNAGA, T., NAYAR, S. K. 1999. Radiometric Self Calibration. *IEEE Comp. Vision & Pattern Recognition*, 1, 374-380.
- NAYAR, S. K., PERI, H., GROSSBERG, M. D., and BELHUMEUR, P.N. 2003. A Projection System with Radiometric Compensation for Screen Imperfections. *ICCV Workshop on Projector-Camera Systems*.
- PERONA, P., and MALIK, J. 1990. Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 7, 629-639.
- RASKAR, R., WELCH, G., LOW, K. L., and BANDYOPADHYAY, D. 2001. Shader Lamps: Animating Real Objects With Image-Based Illumination. In *Proc. of Eurographics Workshop on Rendering Techniques*, 89-102.
- SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., and LENSCH, H. P. A. 2005. Dual Photography. In *Proc. of ACM SIGGRAPH 2005*, 745-755.
- SUN, J., YUAN, L., JIA, J., and SHUM, H. Y. 2005. Image Completion with Structure Propagation. *ACM Trans. on Graphics*, 24, 3, 861-868.
- TOMASI, C., and MANDUCHI, R. 1998. Bilateral Filtering for Gray and Color Images. In *Proc. of IEEE International Conf. on Computer Vision*, 836-846.
- WEI, L. Y., and LEVOY, M. 2000. Fast Texture Synthesis Using Tree-structured Vector Quantization. In *Proc. of ACM SIGGRAPH 2000*, 479-488.
- WEISS, B. 2006. Fast Median and Bilateral Filtering. *ACM Trans. of Graphics*, 25, 3, 519-526.
- WETZSTEIN, G., and BIMBER, O. 2007. Radiometric Compensation through Inverse Light Transport. In *Proc. of Pacific Graphics*, 391-399.