

Projector Placement Planning for High Quality Visualizations on Real-World Colored Objects

Alvin J. Law, Daniel G. Aliaga, *Member, IEEE*, and Aditi Majumder, *Member, IEEE*

Abstract—Many visualization applications benefit from displaying content on real-world objects rather than on a traditional display (e.g., a monitor). This type of visualization display is achieved by projecting precisely controlled illumination from multiple projectors onto the real-world colored objects. For such a task, the placement of the projectors is critical in assuring that the desired visualization is possible. Using ad hoc projector placement may cause some appearances to suffer from color shifting due to insufficient projector light radiance being exposed onto the physical surface. This leads to an incorrect appearance and ultimately to a false and potentially misleading visualization.

In this paper, we present a framework to discover the optimal position and orientation of the projectors for such projection-based visualization displays. An optimal projector placement should be able to achieve the desired visualization with minimal projector light radiance. When determining optimal projector placement, object visibility, surface reflectance properties, and projector-surface distance and orientation need to be considered. We first formalize a theory for appearance editing image formation and construct a constrained linear system of equations that express when a desired novel appearance or visualization is possible given a geometric and surface reflectance model of the physical surface. Then, we show how to apply this constrained system in an adaptive search to efficiently discover the optimal projector placement which achieves the desired appearance. Constraints can be imposed on the maximum radiance allowed by the projectors and the projectors' placement to support specific goals of various visualization applications. We perform several real-world and simulated appearance edits and visualizations to demonstrate the improvement obtained by our discovered projector placement over ad hoc projector placement.

Index Terms — large and high-resolution displays, interaction design, mobile and ubiquitous visualization.

1 INTRODUCTION

Visualization displays are often large and immersive and are mostly composed of simple shapes like planes, piecewise planes (e.g., CAVEs), cylinders, or domes with simple colors (e.g., white and Lambertian). These displays provide a simple surface on which light can be projected to visualize any given data. However, since data for visualizations often originate from or are due to objects with complex shapes and colors (e.g., stress forces on an engine part, weather patterns on a relief map, or deterioration of cultural heritage) using a traditional display system to visualize such data alienates the audience from the reality of the complex physical shape and color of the object itself, thus placing users in a virtual domain of the simple shape and color of the display. In contrast, projecting visualization data directly onto the object's surface enables a richer, more realistic, and potentially more intuitive visualization. The audience can enjoy the benefits of the natural cues of depth perception, parallax, and physical inspection of the real object. In the examples above, causes of stress on an object could be explored in more detail, weather patterns could be correlated with a physical relief map, and deterioration of cultural heritage could be more closely examined. Thus, the ability to produce such a visualization on top of physical objects is an impactful visualization tool.

We call the process of changing the appearance of physical objects *appearance editing visualization* (e.g., [2], [6], [14], [25]).

-
- Alvin J. Law and Daniel G. Aliaga are with Department of Computer Science at Purdue University, IN. Email: {ajlaw|aliaga}@purdue.edu.
 - Aditi Majumder is in the Department of Computer Science at University of California Irvine, CA. Email: majumder@ics.uci.edu.

Manuscript received 31 March 2010; accepted 1 August 2010; posted online 24 October 2010; mailed on 16 October 2010. For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

While augmented reality also superimposes synthetic content on real objects (e.g., [1]), it often needs head-mounted displays or supports a limited number of viewers. In contrast, by carefully controlling how an object is illuminated using digital projectors, we essentially obtain stereoscopic imagery for any number of observers with everything visible to the naked eye without head-mounts or goggles. Despite the setup taking some time, which is insignificant for static scenes (e.g., in a museum), the result is a compelling visualization ability.

1.1 Challenges

There are three main challenges regarding projector placement for appearance editing visualizations on real colored objects. Previous projector placement has generally been ad hoc and informal -- most methods manually attempt to maximize the illumination coverage of the projector on the object surface. The challenges we address are the following:

1. How well the desired visualization is achieved, or its *compensation compliancy*, is reliant on the projectors being able to satisfactorily edit the surface appearance of the objects. Consider a point on an object's surface that has a spectral reflectance curve with small values near the red frequencies (i.e., a small albedo in the red channel). If the desired appearance of the point is a bright red color, then one projector might be insufficient. Such compensation incompliance would result in a color shift, and the resulting visualization could be potentially misleading. To achieve a higher compensation compliancy, more light radiance from additional projectors is required. Since the maximum light radiance from each projector pixel is attenuated by the surface's orientation and distance with respect to the projectors, compensation compliancy is strongly related to projector placement.
2. In addition to the demand for accurate, vivid colors, minimal use of projected light radiance, or *light radiance efficiency*, is also important. This may be due to the demands of the visualization

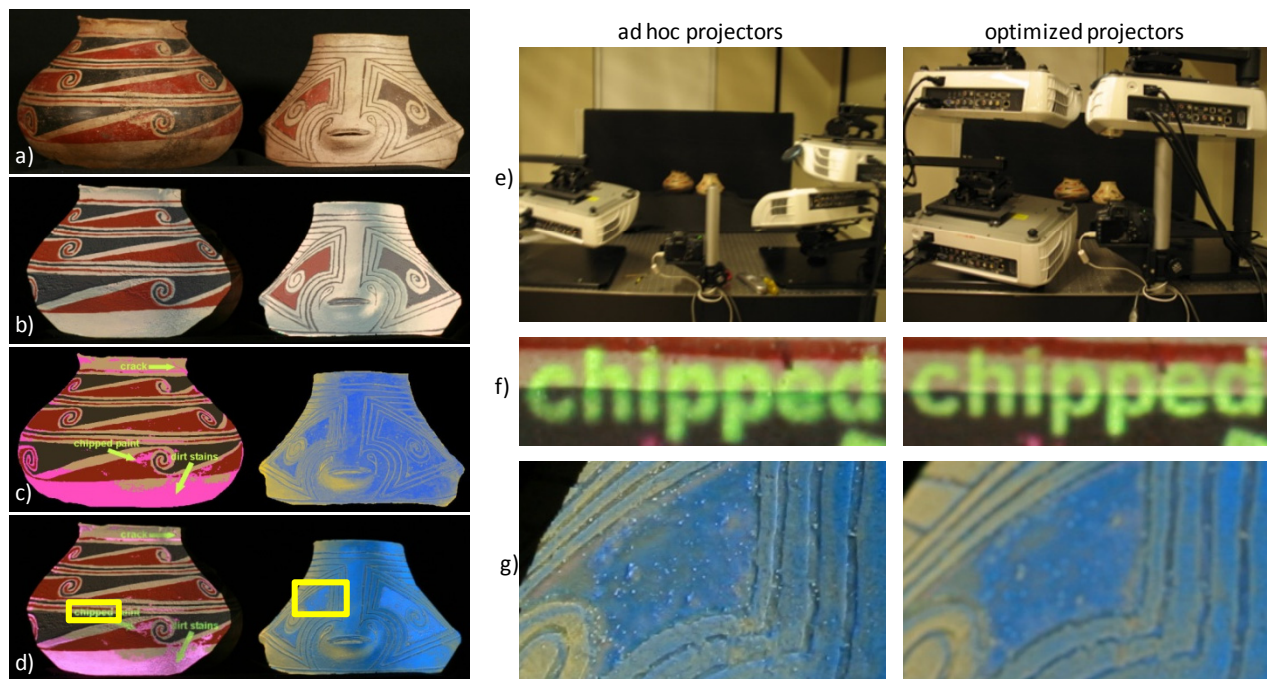


Fig. 1. Restoration visualizations. a) Photograph of the original deteriorated vessels. b) Photograph of the vessels virtually restored using appearance editing. c) Example visualization targets for the vessels (i.e., a synthetic image). d) Photograph of the vessels with (c) as the target visualization appearance. In the left visualization, we restore the appearance of the object but produce a magenta-color appearance to all severely deteriorated surface points (i.e., surface points with the color difference between actual and restored exceeding a threshold). In the right visualization, we use a blue-yellow color map to show the amount of deterioration using a continuous scale (e.g., blue=no deterioration, yellow=significant deterioration). e) The ad hoc and our discovered optimized projector placement used to create (f) and (g). f) A zoom in of the ‘chipped’ annotation on left object. The ad hoc placement results in less crisp yellow text over the background whereas the optimized placement produces sharper text and a more consistent yellow. g) A zoom in of the right object. The ad hoc placement results in a blotchy blue appearance (with some intrinsic red bleeding out). With our optimized placement, the blue is more uniform.

application (e.g., to limit light exposure on cultural heritage when visually restoring them or to accommodate a number of superimposing visualization effects) or limitations of the light sources (e.g., portable projectors, pico-projectors). The amount of light that a projector pixel needs to radiate in order to accurately appearance edit a point on the object depends on the orientation and distance of the projectors relative to that point. Moreover, for colored objects, using more light can increase compensation compliancy while decreasing light radiance efficiency. Therefore, a careful balance must be obtained between compensation compliancy and light radiance efficiency, and this balance can be optimized via careful projector placement.

3. *Practical issues* are also significant. For example, providing a clear path for the user to view the visualization (e.g., avoid blinding them or creating shadows) requires imposing constraints on the projector placement.

1.2 Overview

To the best of our knowledge, we present the first framework that discovers in a holistic manner an optimal projector placement considering the aforementioned three issues. We define an optimal projector placement to be a weighted combination of achieving the desired visualization appearance as best as possible while using as little projector light radiance as possible. This optimality enables our framework to be flexible enough to satisfy a variety of unique visualization applications (Figures 1, 6, 7, and 8). We also show how our framework can be used to cap the amount of light radiance projected per projector pixel and to avoid placing projectors in particular areas to avoid conflict with the observers. Our framework supports imparting visualizations onto diffuse surfaces, and we assume no inter-reflections or other indirect illumination due to the projector light.

Our system calculates the optimal projector positions using an

optimization (Figure 2). A surface model is acquired and the light transport from each projector pixel to the object surface is modeled. Then, given a target visualization and a desired number of projectors, optimal projector position(s) are computed using an adaptive search algorithm. Finally, with the projectors positioned, the surface can be edited.

We present our work in two sections. First, we formalize a theory behind compensation-compliant appearance editing visualizations (Section 3). We define a set of compensation compliancy inequalities which determines whether or not a particular projector combination is capable of achieving a desired visualization on an object. The framework of the inequalities is flexible enough to support multiple color channels (e.g., RGB), to model projector-object visibility, to consider the distance and incidence angle of projector light to the object’s surface, and to support any combination of actual and desired appearances. The inequalities are translated into a set of equalities which are solved with a constrained linear optimization. We also demonstrate how to augment the problem to achieve a solution that balances compensation compliancy with light radiance efficiency.

Second, we provide an efficient search algorithm which applies our constrained linear system to actual placement problems (Section 4). Through the use of a sample-based adaptive projector placement approach and a sparse sampling of the object’s surface points, we efficiently discover an optimal projector placement for a desired balance of compensation compliancy and light radiance efficiency. Our method follows a heuristic of broadly sampling a range of allowable projector combinations, identifying the projector combinations which tend to yield the most optimal appearance editing visualizations, and adaptively narrowing down the search region. The sparse sampling is designed carefully to achieve computational acceleration without compromising the quality of the appearance editing visualization. Finally, we show how this search can be guided to

converge to a solution that assures practical placement constraints.

To demonstrate our approach, we perform several experiments in simulation and on real colored objects. We compare the appearances obtained using optimized projector placement to the appearances obtained using ad hoc projector placement. We show visualization applications where maximizing light radiance efficiency plays an important role and demonstrate imposing practical placement constraints.

2 RELATED WORK

2.1 Displays for Visualization

A large body of literature exists on creating immersive projection-based displays for visualization. These works address calibrating projectors geometrically on planar surfaces (e.g., [4], [5], [8]), on cylindrical surfaces (e.g., [27], [28]), on arbitrary surfaces [25]; and on achieving color uniformity across multiple projectors (e.g., [19], [20], [26]). However, in all these works the display shape is relatively simple, even if immersive. In contrast, we are the first to address projector placement planning when considering the arbitrary shape of the real colored object itself as the display and visualizing data on the physical object itself. Further, unlike the careful projector placement in our method, all the aforementioned techniques casually place multiple projectors in a tiled manner without planning their placements in any way at all.

2.2 Appearance Editing

The ability to produce desired visual content on top of an arbitrary physical surface or set of objects is an important function of many visualization and computer graphics applications, including augmented reality, everywhere displays, telepresence, and telecollaboration (e.g., [1], [14], [16], [17], [18], [24]). Most work in appearance editing has focused on computing the compensation image (i.e., colors to project) in order to alter an object’s appearance (e.g., [3], [7], [22], [25]). In these works, projectors are placed in ad hoc positions which only seem intuitively proper. Further, these works mostly deal with near-white objects (i.e., a high spectral reflectance curve for all visible wavelengths) thereby reducing the need for projector planning. Hence, projector placement has generally not been addressed in appearance editing.

Changing the appearance of colored objects is more difficult since the intrinsic object colors have to be neutralized before imparting a new appearance onto the object’s surface. While methods handling such radiometric compensation have been explored (e.g., [14], [16], [21], [32]), these methods only use a single projector to achieve appearance changes with planar or smoothly curved surfaces – placement of a single projector in these cases is generally straightforward. In Aliaga et al. [2], multiple projectors are used for virtual restoration of deteriorated and colored artifacts. Their method pursues luminosity compliancy in the sense of attempting to ensure a desired new appearance only seeks to make points on the object appear darker relative to the original appearance. However, per-channel compliancy, light radiance efficiency, and projector placement are not considered. Damera-Venkata and Chang [11] use multiple superimposed projectors to achieve display super-sampling but assume near-planar surfaces and do not consider appearance quality improvements that may result from altering projector positions. With our framework, the range of possible appearance edits and the light-radiance efficiency of the solution is improved.

2.3 Viewpoint Planning

Viewpoint planning is a large field of study that overlaps with computer graphics, computer vision, and robotics. The visibility problem of determining an optimal set of viewpoints from which most or all of the object’s surfaces are visible is NP hard and is related to the art

```

1. Acquire 3D model of scene.
2. Define desired target appearance.
3. Set initial angular range and sampling
   numbers for projector positions.
4. Select desired balance of compliancy
   and light radiance efficiency.
5. WHILE (Q metric improving)
   a. FOR all unique projector combos
       Compute visibility term.
       Solve  $Mx=t$ .
     ENDFOR
   b. Normalize compensation compliancy
       and light radiance eff. values.
   c. Sort all solutions using normalized
       metric Q.
   d. Pick best solution, re-center and
       subdivide angular ranges.
   ENDWHILE
6. Setup projectors and start appearance
   editing application.

```

Fig. 2. Adaptive Placement Algorithm. We show a pseudo-code summary of our placement algorithm.

gallery problem [23] and to the aspect graph [15]. Heuristic approaches have been devised for 3D acquisition (see [30] for a survey) and for synthetic rendering (e.g., [12], [33]). These methods focus on surface coverage, from-point/from-region visibility, and, in some cases, on ensuring that visibility at grazing-angles is reduced. However, in order to improve altering an object’s appearance with projected light, we analyze the interaction between the per-color channel surface albedo and the light radiance. Moreover, we also address controlling light-radiance efficiency and imposing spatial constraints.

3 COMPENSATION COMPLIANCY FORMULATION

Multiple factors must be considered when designing a formulation to determine whether a target appearance or visualization is possible. The formulation should consider (i) the amount of light radiance each projector can contribute to any one point on the object’s surface based on the projector-point visibility, (ii) the angle between the outgoing vector to each projector’s position and the surface point’s normal, (iii) the surface point’s albedo, and (iv) the attenuation due to the distance from the projector to the surface point.

In the subsections that follow, we describe appearance editing theory in terms of a single intensity channel for clarity. These equations can be generalized to any number of channels (e.g., RGB) by simply repeating the equations for each color channel. Further, without loss of generality we assume all projectors to be of similar maximum luminance.

3.1 Compensation Compliance

To achieve a compensation-compliant appearance, all of the points on the surface must be able to change the color resulting from their true albedo to a color that approximates the target appearance. Once sufficient color radiance can be reflected back to the observer, any additional light potentially incident on a surface point is superfluous and will not further improve the compliancy of the surface point. We therefore model compensation compliance using an inequality of light contribution on the surface attenuated by the various aforementioned factors of surface-projector interaction.

The *compensation-compliance inequality* for a surface represented by points q_i , for $i \in [1, N]$, in a setup with projectors at positions p_j , for $j \in [1, P]$ (Figure 3), is expressed as:

$$\sum_{j=1}^P (A(q_i, p_j) + b_j) + b_a - T_i \geq 0 \quad (1)$$

where $A(q_i, p_j)$ represents the function generating the appearance possible at object location q_i using a projector at position p_j and T_i represents the function calculating the target appearance at location q_i . We do not optimize for the direction of projection; rather, we assume the projector is pointing towards the object/surface and has sufficient field-of-view to illuminate the object/surface. The constant b_j is the minimum black-level achievable by a projector j , and b_a is the black level ambient lighting in the assumed dark environment. All b_j 's and b_a are disregarded during projector placement planning because 1) we assume these values are small and thus negligible, and 2) the addition of constants does not impact the relative values of the Q metric (discussed in Section 3.3) during optimization. An appearance is compensation compliant if all surface points' inequalities are simultaneously satisfied.

We define an explicit representation for the appearance generating function A by assuming a Lambertian image formation process on the surface of the object similar to those used in prior appearance editing works (e.g., [2], [16], [25]). Using surface normals n_i and albedos α_i , a maximum light radiance I_{max} for each projector, attenuation inversely proportional to the square of projector-point distance, and a binary visibility term (i.e., a point q_i is either visible or not by p_j), the function A can be expanded as

$$A(q_i, p_j) = \left(\frac{1}{\|p_j - q_i\|^2} \right) I_{max} v_{ij} \alpha_i \left(n_i \cdot \left(\frac{p_j - q_i}{\|p_j - q_i\|} \right) \right) \quad (2)$$

The function T , which represents the target appearance, is defined by the visualization application – it need not follow the model used for the appearance generation. In our experiments, we create target appearances by using an acquired 3D model of the object and Gouraud or fixed-viewpoint Phong shading, optionally with texture mapping. Although we could write T_i as a function in terms of q_i , n_i , a desired albedo α'_i , and the synthetic light parameters, for brevity we simply let T_i be the target color for the projection of object point q_i onto the camera image.

3.2 Constrained Linear System

To convert the nonlinear compensation compliance inequality into a linear system of equations, we insert additional unknowns borrowed from radiometric compensation calculations and make a few simplifying assumptions. Combining equations (1) and (2) yields a nonlinear system of inequalities with each surface point contributing an inequality. Naïvely using these inequalities to determine a projector placement permits many configurations as potential solutions. However, the solutions vary significantly in quality. For example, a placement solution can position projectors at a grazing angle thus




Fig. 3. Projector Setup. A projector at p_j is adaptively computed by subdividing the sampling range of spherical coordinates θ_j and ϕ_j . Visibility of surface point q_i from p_j is represented by v_{ij} . Surface albedo α_i , distance and orientation of the projector from the surface, and maximum projector radiance determine if a target visualization is possible.

requiring the projectors to operate at or near maximum radiance in order to achieve the desired colors. Such a placement may also decrease the apparent resolution and sharpness of the appearance that can be accomplished by the projectors – this is because the footprint of each projector pixel on the object's surface is much larger at a grazing angle. Also, the inequality of equation (1) implicitly assumes each projector pixel is operating at maximum radiance. In actuality, the radiance of each projector pixel is controllable and its value is usually computed during a radiometric compensation process (e.g., [16]). Thus, we need to alter the naïve usage of equations (1) and (2) into a linear system of equations that permits controlling the quality of the solution while also being efficient to compute.

3.2.1 From Inequality to Equality

To convert the compensation compliance inequality into an equation, we insert per-object point intensity values $s_{ij} \in [0,1]$ into equations (1) and (2). The range constraining is needed because each projector pixel can only radiate light from its constant black-level b_j (assumed to be zero) to a maximum of I_{max} . After algebraic manipulations we obtain

$$\sum_j^P \left(\frac{I_{max} v_{ij} s_{ij} \alpha_i (n_i \cdot (p_j - q_i))}{\|p_j - q_i\|^3} \right) = T_i \quad (3)$$

which is still nonlinear because the unknowns (v_{ij} , s_{ij} , and p_j) are multiplied and p_j appears in the denominator.

3.2.2 Linearity

To convert equation (3) into a linear system of equations, we make some simplifications that will leave s_{ij} as the only unknown. First, we make each distance term in the denominator of (3) a constant by restricting projector positions to being on a subset of the surface of an imaginary sphere surrounding the scene. This restriction allows us to pre-compute each projector-point distance prior to each solving of the linear system. While this restriction fixes the distance from the projectors to the scene's center, the scene can still consist of multiple objects at varying distances to the projectors.

Second, to make s_{ij} the only unknown in the numerator of equation (3), we make the visibility term v_{ij} a constant by setting it to one of a set sampled values. Accounting for visibility (e.g., self-occlusion of the target object, projector-object occlusions, etc.) is a challenging problem encountered in many geometric algorithms. Rather than finding a closed-form representation, we compute a set of projector position combinations and evaluate the visibility in each case to yield binary values for v_{ij} . We sample the subset of the spherical surface on which each projector can be placed with $S = S_\theta \times S_\phi$ positions, where S_θ and S_ϕ are the number of samples for the θ and ϕ spherical coordinates. For exactly P projectors, where at most one projector can occupy a given position, we sample $\binom{S}{P} = \frac{S!}{P!(S-P)!}$ configurations. Thus, the total number of configurations, using from 1 to P projectors, is $C = \sum_{j=1}^P \frac{S!}{j!(S-j)!}$. Section 4 will describe how we adaptively perform visibility sampling and select placement configurations.

3.2.3 Appearance Resolution

For our goal of determining compensation compliancy, highly accurate projector pixel modeling is not crucial, but we do want to capture the resolution (i.e., detail level) present in the target appearance. The footprint of a projector pixel on the object's surface appears on the camera image as a finite-sized patch rather than an infinitesimal point. To model this mapping of projector pixels to camera pixels, several methods have been proposed (e.g., Gaussian elliptical pixels [10] and weighted elliptical points [34]). In our system, the camera has roughly 10 times more pixels than the projector and both the

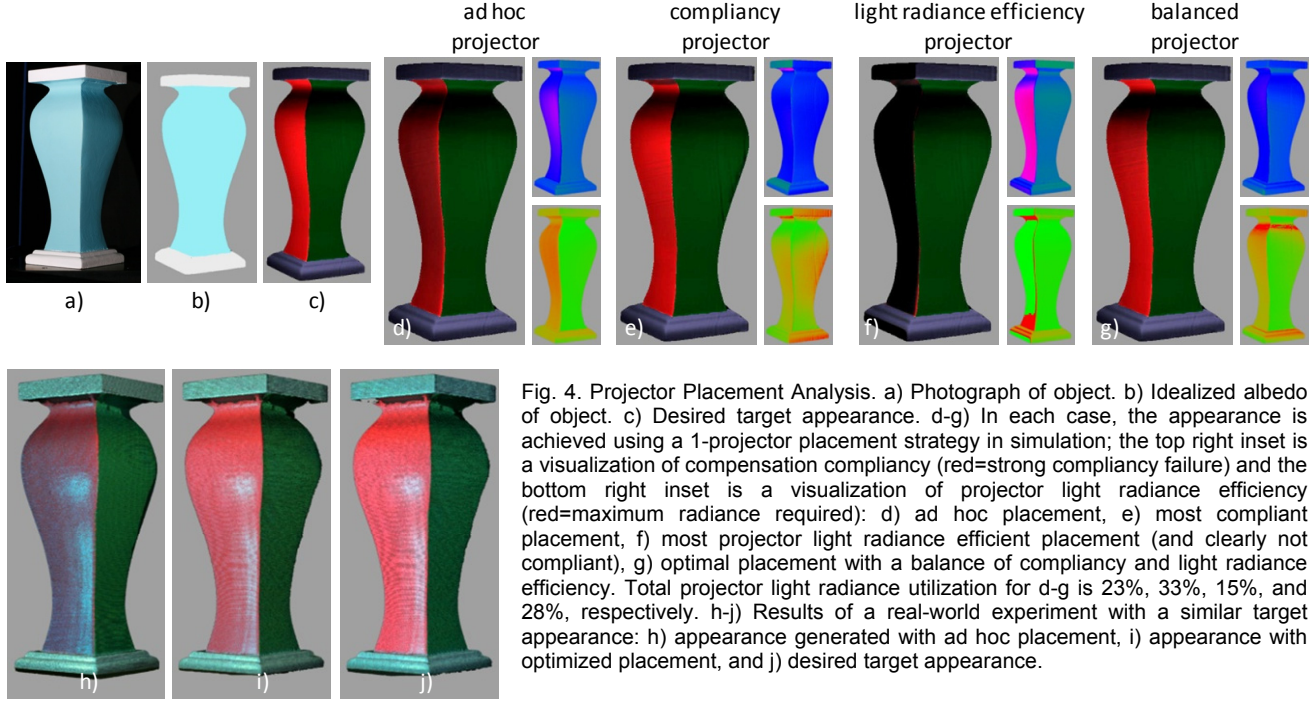


Fig. 4. Projector Placement Analysis. a) Photograph of object. b) Idealized albedo of object. c) Desired target appearance. d-g) In each case, the appearance is achieved using a 1-projector placement strategy in simulation; the top right inset is a visualization of compensation compliancy (red=strong compliancy failure) and the bottom right inset is a visualization of projector light radiance efficiency (red=maximum radiance required): d) ad hoc placement, e) most compliant placement, f) most projector light radiance efficient placement (and clearly not compliant), g) optimal placement with a balance of compliancy and light radiance efficiency. Total projector light radiance utilization for d-g is 23%, 33%, 15%, and 28%, respectively. h-j) Results of a real-world experiment with a similar target appearance: h) appearance generated with ad hoc placement, i) appearance with optimized placement, and j) desired target appearance.

projectors and the cameras are at a similar distance to the scene. Thus, we model the footprint of a projector pixel on the camera image as a small patch of 3x3 or 5x5 camera pixels. We assume the same s_{ij} value is used for all object points in the same patch (i.e., they are all being illuminated by s_{ij} 's projected radiance). Since the original object points will be sparsely sampled (as described in Section 4), the patches remain disjoint, preventing a situation where multiple pixels from multiple projectors illuminate the same surface point resulting in s_{ij} 's being interdependent.

Altogether, the new linear equation for compensation compliancy takes the form

$$M_k x_k = t_k \quad (4)$$

where

$$M_k = I_{max} \begin{bmatrix} W_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & W_{N_k} \end{bmatrix} \quad W_i = \left[\frac{\sum_j \alpha_i(n_i \cdot d_{ij})}{\|d_{ij}\|^2} \quad \dots \right]^T$$

$$x_k = [s_1 \dots s_{N_k}]$$

$$t_k = [T_1 \dots T_{GN_k}]$$

and $k \in [1, C]$ represents the current sampled visibility case, N_k is the number of object points visible in the current sampled visibility case, G is the number of objects points in each patch (e.g., 9 or 25), W_i is a column vector representing the maximum possible intensity value attainable by the P projectors for each of the G points of a patch (each point corresponding to different target values in t_k), $d_{ij} = p_j - q_i$ (pre-computed), and the solution x_k is solved while subject to the linear constraint $0 \leq x_k[m] \leq 1$ for $m \in [1, N_k]$ (e.g., using [8]). Since only points visible by a projector are used, $N_k < N$.

3.3 Compliancy and Projector Light Radiance

To obtain a light radiance efficient projector placement, we desire the components of the solution vector to be small. If M_k in equation (4) is constructed for a scene assumed to have no indirect illumination effects, then M_k is of full rank and the corresponding solution vector x_k is unique. However, even in this simple case, there might be other equally compliant (or slightly less compliant) projector configurations with solutions requiring less overall per projector pixel radiance (e.g., $\|x_{k'}\| < \|x_k\|$ for some $k' \in [1, C]$ and $k' \neq k$), and such a solution might be preferred; e.g., at half the light radiance, we might

obtain a solution that is only 10% less compliant. This defines a tradeoff between full compliancy and best projector light radiance efficiency.

We use a metric Q_k that, given a user-defined relative-importance value, computes a weighted sum of compliancy and light radiance efficiency. High-compliancy might be desired when visualization quality is of the utmost importance. Light radiance efficiency might be needed in order to reduce the projected radiance required and/or the number of projectors needed. Discovering the smallest number of required projectors can be done by increasing the number of projectors until the thresholds are met.

One definition for the metric Q is defined as

$$Q_k(\beta) = \frac{\beta(c_k + \sum_{i=1}^N T_i \delta_i)}{C_{max}} + \frac{(1-\beta)l_k}{L_{max}} \quad (5)$$

where $c_k = \|x_k - M_k^+ t_k\|$ and $l_k = \|x_k\|$ are the inverse compensation compliancy and the light radiance efficiency measures, M^+ is the pseudo-inverse of M , C_{max} and L_{max} are constants to normalize each respective equation term to the range $[0, 1]$, and user-provided $\beta \in [0, 1]$ indicates the relative importance of compensation compliancy and light radiance efficiency. The summation of T_i 's within the numerator in equation (5) accounts for the surface points not visible by any projector to ensure high surface coverage by the projectors (e.g., a solution of “nothing visible” does not yield a compliant compensation). The variable δ_i is a Kronecker delta with a value of 1 when object point q_i is not visible by any projector and 0 otherwise. Another definition for Q is the ratio of light radiance efficiency (cost) to the compliancy sum (benefit). We explore both definitions in our results section.

4 PROJECTOR PLACEMENT ALGORITHM

Given our formulation for compensation compliancy, we efficiently generate an optimal projector placement combination that results in high-quality appearance editing visualizations. To yield an efficient algorithm, we (i) use an adaptive sampling technique to discover the optimal P -projector combination and (ii) perform a sparse sampling of the object's surface points, used in our linear system of equations, that – despite the sparsity of the sampling – attempts to achieve the visual resolution implied by the target appearance.

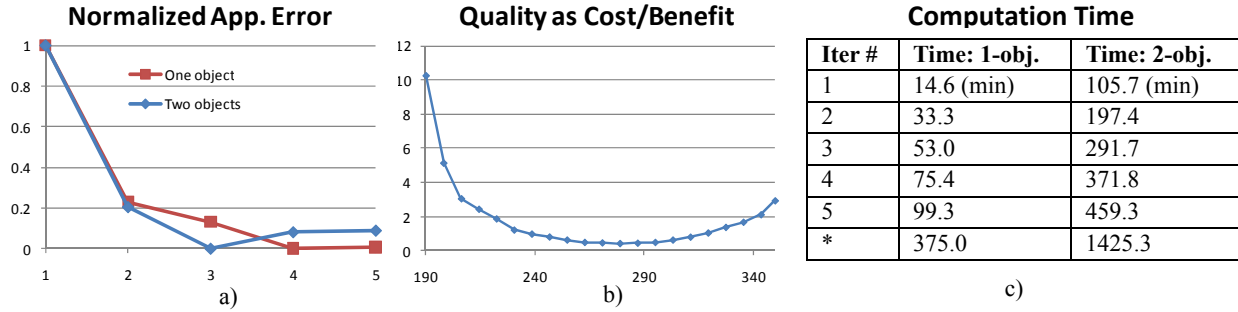


Fig. 5. Behavior of Adaptive Algorithm. a) We show the change of the normalized appearance error during successive adaptive iterations of our algorithm for two scenes. b) We demonstrate the cost/benefit ratio value resulting from changing the ϕ coordinate of a projector placed in front of an example scene. c) A representative summary of the computation times of our algorithm. Times are for those graphed in (a) which start at θ and ϕ using 4 and 6 samples spanning a hemisphere with the angular range being halved each iteration. The last row is the time needed for a dense sampling of 6 and 16 samples for θ and ϕ . It yields an error almost identical to the final error of the adaptive approach.

4.1 Adaptive Placement

Our adaptive sampling method steers the computational cost towards projector combinations which yield more promise in discovering an optimal projector placement combination. Finding the best multi-projector placement by a brute force sampling of all P -projector placement combinations and their object-projector visibility is impractical. Instead, our method broadly samples the range of possible projector combinations and identifies the projector combination which yields the best balance of compensation compliancy and light radiance efficiency. Then, the method recursively narrows down the range of the projector placement parameters centered on the best parameter values.

In our method, the user defines initial ranges for the spherical coordinates θ and ϕ for each of P projectors, a subdivision factor f to narrow the ranges for θ and ϕ per iteration, and a maximum number of iterations. After each iteration, we use our metric Q to evaluate the quality of all sampled projector placements. There could be multiple clusters of high-quality configurations within the $2P$ dimensional space sampled in each iteration (i.e., there are P projectors and each projector is represented by a sampling of the tuple $[\theta, \phi]$). However, in our experiments we observe that Q changes in a locally smooth manner over the sampled positions; thus, we naively subdivide the projector placement ranges surrounding the best discovered case. For scenes with frequent and abrupt changes in visibility, Q may vary more abruptly. In these situations, one may capture this behavior either by sampling θ and ϕ more densely or by subdividing into multiple projector placement ranges per iteration.

4.2 Sampling

Since our acquired models of the target objects are high resolution (e.g., 10^5 , 10^6 , or more vertices), it is impractical to use every vertex/point in each use of equation (4). Instead, we sparsely sample the vertices for a representative set of points. We found 100-500 randomly distributed points over the potentially visible surface of each object to be sufficient.

4.3 Q-Metric Normalization

Evaluating Q requires normalization of compensation compliancy and of light radiance. The values for C_{max} and L_{max} are not known a priori. Hence, we re-estimate them during each iteration. Using all of the solutions to equation (4) for the current iteration (i.e., x_k for $k \in [1, C]$), we compute the c_k 's and l_k 's and their means \bar{c} and \bar{l} , standard deviations σ_c and σ_l , minimums m_c and m_l , and maximums M_c and M_l . Then, we ignore mostly non-compliant solutions (i.e., $c_k > \bar{c} + 2\sigma_c$) and ignore excessively high light radiance solutions (i.e., $l_k > \bar{l} + 2\sigma_l$). We update the minimum and maximum ranges

for each of the two norms using the remaining solutions and update the norms to be $c'_k = (c_k - m_c)$ and $l'_k = (l_k - m_l)$. Lastly, we define $C_{max} = M_c - m_c$ and $L_{max} = M_l - m_l$. This dynamic per-iteration normalization procedure helps to maintain the desired balance between compensation compliancy and light radiance efficiency.

For multi-projector configurations, some symmetric projector placement arrangements can occur. One case of symmetry is caused by the identical projector assumption; we account for this when sampling projector placement configurations (Section 3.2.2). Another case of symmetry is produced by scene geometry: similar illumination scenarios can be caused by symmetries present in the scene geometry. We do not explicitly account for these rare occurrences.

5 RESULTS AND DISCUSSION

We have used our placement algorithm in several simulated and real-world experiments. Our experimental system uses a self-calibrating

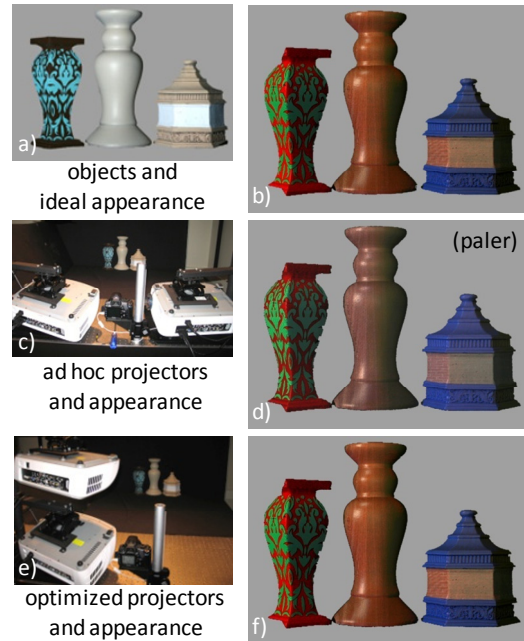


Fig. 6. Multiple Objects and Projectors. a) Photograph of the objects. b) Ideal appearance. c-d) Ad hoc projector placement yields an incorrect, paler appearance. e-f) Optimized projector placement produces an appearance which matches the ideal appearance. Some geometry is clipped due to lack of visibility to at least one of the projectors.

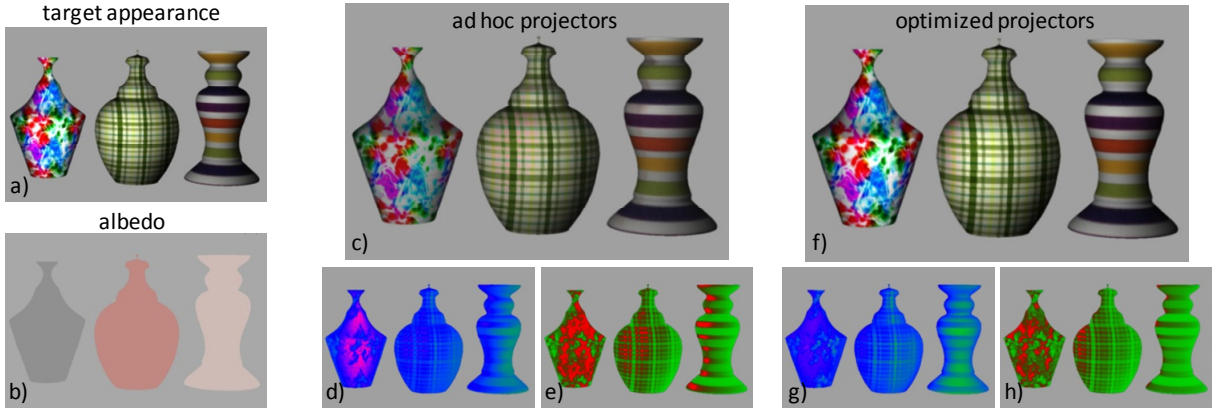


Fig. 7. Projector Placement Planning. A rendering of a synthetic appearance edited scene of 3 objects with target appearance (a) and surface albedo (b). c) Using two simultaneous projectors placed ad hoc in front of the scene and along a semi-circle, the resulting appearance is paler. d) A visualization of the compliancy achieved and e) a visualization of the amount of projector light radiance needed to produce (c). f) A rendering of the same scene but using optimized projector placement computed by our method (projectors biased towards the left object). The obtained imagery is noticeably brighter, of higher contrast, and more similar to the target appearance shown in (a). g) The compliancy visualization indicates the target appearance is better achieved and h) the projector light radiance visualization indicates a decrease in the total amount of projector light needed.

structured light system (e.g., [12], [31]) to digitize the scene. We perform a radiometric calibration of the projectors/objects using Nayar et al. [21]. We generate visualizations by using the 3D captured model and a custom program exploiting OpenGL and texture-mapping. Since we have a mapping from camera pixels to the 3D model, we know the desired target color for each object point (i.e., the values for the t_k vector). Our simulated experiments use the geometry and idealized albedos of actual objects that have been 3D acquired. Further, the projectors are virtually placed at the optimized locations. In contrast, for a real-world experiment, the albedo is estimated via photometric stereo and projectors are manually placed near their optimal positions.

5.1 Projector Placement Analysis

Figure 4 shows an analysis of the effects of altering the placement of a single projector. The object is shown in Figure 4a. The idealized albedo values (i.e., albedo values manually established for the acquired 3D object without any shading effects) are shown in Figure 4b. Figure 4c shows the desired target appearance – strong changes to the color of the object are pursued. Figures 4d-g show the synthetic result of various placements of the projector. Along with the resulting appearance, we show a compliancy map (green indicates the target colors are easily achieved, blue implies compliancy with little excess projector radiance, and red indicates strong lack of compliancy) and a light radiance efficiency map (green implies low light radiance used and red implies high light radiance used). Figure 4d visualizes the compliancy and light radiance efficiency obtained by an ad hoc projector directly in front of the object. Figure 4e shows the results when the projector is placed at the computed location for best compliancy (i.e., $\beta = 1$). While the target appearance is achieved, significantly more light than the other configurations is needed. Conversely, Figure 4f contains the visualizations for placing the projector at the location for best light radiance efficiency (i.e., $\beta = 0$). This configuration clearly is not able to produce the target appearance. Lastly, Figure 4g shows the results obtained when using the projector placement computed by our algorithm for a balance of compliancy and light radiance efficiency (i.e., $\beta = 0.5$). In this example, the optimized placement obtains compliancy similar to Figure 4e while using less projector light.

Figures 4h-j contain photographs from a real-world experiment using the same object and a similar desired appearance. Figure 4h shows a photograph of the appearance compensated object using the

ad hoc placement, Figure 4i shows a photograph of the appearance modified object using the optimal placement (biased towards the red surface), and Figure 4j depicts the desired target appearance generated using an unbounded amount of projector light. All three images show an overall color change from synthetic because they are captured with a digital camera and use radiometric calibration. While the maximum projector light radiance is the same for both the ad hoc and optimal projector placements, the optimal placement shows higher brightness and contrast due to better use of the available projector light.

5.2 Time and Error Behavior of Adaptive Algorithm

Figure 5 depicts the behavior of the adaptive algorithm. Figure 5a shows how the appearance error (i.e., inverse compliancy measure c_k) changes after each iteration. The graph contains one curve for the object shown in Figure 4 and one curve for the two-vase, two projector synthetic scene (shown at end of accompanying video; omitted from paper because of space). The best compliancy value (i.e., smallest c_k) is mapped to zero and the compliancy after the first iteration is mapped to one to normalize the curves. Both curves show a clear improvement with iterations although not strictly monotonic – this is due to sampling that might place the projector at a (near) optimal placement by chance. Usually 4 to 5 iterations are sufficient to reach an optimal placement.

Our metric Q is typically implemented using equation (5), and the effect of changing β is shown in Figure 4. In Figure 5b, we also show the value of a cost/benefit ratio where cost is required light radiance and benefit is compliancy (i.e., $\text{cost/benefit} = 1/(c_k l_k)$). The graph shows the ratio for the object in Figure 4 as the projector position changes along a semi-circle at $\theta = 0^\circ$ and $\phi \in [190^\circ, 350^\circ]$. In this case, the graph indicates a clear optimal position near 290° .

Figure 5c provides a representative summary of computation times (in minutes). The times shown indicate the amount of time elapsed after each iteration is completed. Our adaptive approach yields a compliancy level similar to that possible using a single iteration of dense sampling after four iterations while using 5x and 4x less computation time. Furthermore, the accuracy of the placement obtained by our repeated subdivision of steps is also better by about 5x (e.g., 5x and 8x in θ and 3x and 3.5x in ϕ). The two-object scene requires more time because more object points are used thus making M_k larger. It is worth noting that our current implementation runs

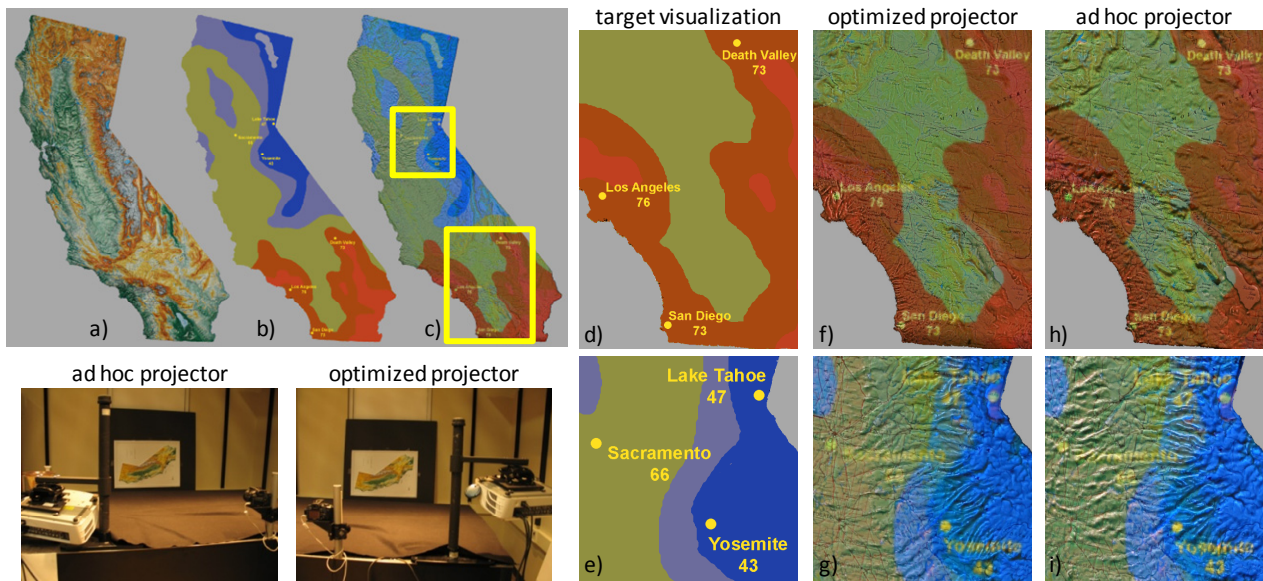


Fig. 8: California Relief Map. A temperature map visualization appearance edited on a physical relief map of California. a) Photograph of the relief map. b) Target visualization appearance. c) Photograph of relief map with visualization. d-e) Zoom-in of the target visualization appearance for two regions of California. f-i) Comparison of the appearance edited visualization for an optimized (f, g) and ad hoc projector placement (h, i). Both projector placements, shown in the bottom left, avoided placing the projector directly in the middle of the setup to prevent obstructing the audience. However, our solution achieves a more accurate visualization in comparing (f) to (h) without losing visual quality comparing (g) to (i).

sequentially on a single CPU core. However, our placement algorithm is highly parallelizable, thus we could significantly reduce the computation time using multiple cores or a GPU-based acceleration scheme.

5.3 Examples using Multiple Objects & Projectors

Figures 1, 6, 7, and 8 show several example scenes using multiple objects and projectors, both in simulation and in real-world experiments. Figure 1 shows a visualization application for the restoration of cultural heritage artifacts. We produce two visualization examples when restoring the age-deteriorated vessels using projected light. On the left of Figures 1a-d, we visually restore the artifact but also draw in magenta-color over all severely deteriorated surface points. On the right of Figures 1a-d, we use a blue-yellow color map to show the amount of deterioration using a continuous scale. Figures 1e-g show the setup and various close-ups of the obtained appearance for both the ad hoc and the optimized placement of three projectors. The placement computed by our method shows a clearly better quality visual appearance.

Figure 6 contains photographs of a real-world experiment using three objects and two projectors. Figures 6a-b show the objects and an ideal appearance. Figures 6c-d shows the appearance possible using an ad hoc projector placement. The resulting appearance is paler than the ideal appearance. Figures 6e-f shows the improved quality obtained by using the optimal projector placement. The resulting image matches the ideal image well.

Figure 7 shows a synthetic scene illuminated by two projectors. Our optimal placement of the projectors enables us to obtain superior brightness and contrast to better match the desired target appearance while using less total projector light than an ad hoc placement (where projectors are placed symmetrically in front of the scene). In the optimized placement, the projectors are carefully positioned more to the left so as to enable altering the more challenging object on the left while not sacrificing the compliancy of the object on the right. A balance of compliancy and light radiance efficiency is also achieved. Although it is demonstrated in this case, a highly-compliant projector placement does not necessarily reduce the maximum needed projector radiance, but if such is possible, our algorithm will

discover it.

Finally, Figure 8 shows a relief map of California with a temperature map visualization. Figures 8a-c show the relief map, target visualization, and resulting visualization from an optimally placed projector directly in front of the map. The colors of the visualization represent different temperatures, and various locations are annotated with their temperatures. For this experiment, we wish to avoid placing the projector in front of the relief map in order to avoid blocking the audience’s view. Naively, one may place the projector to the left of the map (our ad hoc position), but our optimization indicates that placing the projector on the right of the map yields a better visualization – determining the placement is difficult to do by mere inspection of the object and the desired visualization. Figures 8d-e show a zoom in of two areas of the target visualization in Figure 8b. Figures 8f-g show the (better) resulting visualization using the optimized projector placement, and Figures 8h-i show the (worse) resulting visualization using ad hoc placement.

6 CONCLUSIONS AND FUTURE WORK

We have presented a framework to discover optimal projector positions for appearance editing visualization. Projecting visualization data directly onto the object’s surface enables a richer, more realistic, and potentially more intuitive visualization; we look forward to its proliferation. In this work, we pay careful attention to the placement of the projectors and as a consequence can produce high-quality visualizations with less unwanted color distortions and visual artifacts. Our results, both simulated and in the real world, show examples where optimal projector placements achieve what ad hoc projector placements cannot. Further, the optimal placements can achieve the target visualization using a minimal amount of light radiance.

With regards to limitations and future work, there are several items. Our approach does not account for indirect illumination effects. Further, the planning process may produce projector placements that are hard to make use of in practice. In addition, our approach does not scale well to a very large number of projectors but such is typically not the case. For specific future work items, first we would like to explore how to best simultaneously use projectors of

significantly different radiometric behavior (e.g., of different maximum intensity, color response, etc.). Second, we would like to include resolution into the definition of optimal so as to also arrive at the configuration able to produce the highest overall resolution. Third, we would like to improve our algorithm performance by using multiple cores and/or GPU programming. Finally, we would also like to explore including color perception phenomena (e.g., simultaneous color contrast, chromatic adaptation) into planning so as to make the visualization even more compelling.

REFERENCES

- [1] R. Azuma, Y. Baillet, R. Behringer, S. Feiner, S. Juliers, and B. MacIntyre, "Recent Advances in Augmented Reality", *IEEE Computer Graphics and Applications*, 21:6, 34-47, 2001.
- [2] D. Aliaga, A. Law, and Y. Yeung, "A Virtual Restoration Stage for Real-World Objects", *ACM Trans. on Graphics*, 27:5, 2008.
- [3] D. Bandyopadhyay, R. Raskar, and H. Fuchs, "Dynamic Shader Lamps: Painting on Movable Objects", *Proc. of IEEE/ACM Intl. Symposium on Augmented Reality*, 207-216, 2001.
- [4] E. Bhasker, R. Juang, and A. Majumder, "Registration Techniques for Using Imperfect and Partially Calibrated Devices in Planar Multi-Projector Displays", *IEEE Trans. On Visualization and Computer Graphics*, 2007.
- [5] E. Bhasker, P. Sinha, and A. Majumder, "Asynchronous Distributed Calibration for Scalable Reconfigurable Multi-Projector Displays", *IEEE Trans. On Visualization and Computer Graphics*, 2006.
- [6] O. Bimber, B. Fröhlich, D. Schmalstieg, and L.M. Encarnação, "The Virtual Showcase", *IEEE Computer Graphics & Applications*, 21:6, pp. 48-55, 2001.
- [7] O. Bimber and D. Iwai, "Superimposing Dynamic Range", *ACM Trans. on Graphics*, 27:5, 2008.
- [8] H. Chen, R. Sukthankar, G. Wallace, and K. Li, "Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees", *Proc. of IEEE Visualization*, 2002.
- [9] T.F. Coleman and Y. Li, "A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on Some of the Variables", *SIAM Journal on Optimization*, 6(4), 1040-1058, 1996.
- [10] Y. Chuang, D. Zongker, J. Hindorff, B. Curless, D. Salesin, and R. Szeliski, "Environment Matting Extensions: Towards Higher Accuracy and Real-Time Capture", *Proc. of ACM SIGGRAPH*, 121-130, 2000.
- [11] N. Damera-Venkata and N. Chang, "Display Supersampling", *ACM Trans. on Graphics*, 28:1, 2009.
- [12] P. Debevec, Y. Yu, and G. Borshukov, "Efficient View-Dependent Image-Based Rendering with Projective Texture-Mapping", *Eurographics Rendering Workshop*, 105-116, 1998.
- [13] R. Furukawa and H. Kawasaki, "Uncalibrated Multiple Image Stereo System with Arbitrarily Movable Camera and Projector for Wide Range Scanning", *Proc. of 3DIM*, 302-309, 2005.
- [14] A. Grundhöfer and O. Bimber, "Real-time Adaptive Radiometric Compensation", *IEEE Trans. on Visualization and Computer Graphics*, 14:1, 97-108, 2008.
- [15] Z. Gigus and J. Malik, "Computing the aspect graph for line drawings of polyhedral objects", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12:2, 113-122, 1990.
- [16] M.D. Grossberg, H. Peri, S.K. Nayar, and P.N. Belhumeur, "Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System", *IEEE Conf. on Computer Vision and Pattern Recognition*, 1:452-459, 2004.
- [17] M. Gross, S. Würmlin, M. Naef, E. Lamboray, C. Spagno, A. Kunz, E. Koller-Meier, T. Svoboda, L. Van Gool, S. Lang, K. Strehlke, A. Vande Moere, and O. Staadt, "blue-c: A Spatially Immersive Display and 3D Video Portal for Telepresence", *Proc. of ACM SIGGRAPH*, 819-827, 2003.
- [18] T. Johnson and H. Fuchs, "A Unified Multi-Surface, Multi-Resolution Workspace with Camera-Based Scanning and Projector-Based Illumination", *Eurographics Symp. on Virtual Environments-Immersive Projection Technology Workshop*, 2007.
- [19] A. Majumder, Z. He, H. Towles, and G. Welch, "Achieving Color Uniformity Across Multi-Projector Displays", *Proc. Of IEEE Visualization*, 2000.
- [20] A. Majumder and R. Stevens, "Perceptual Photometric Seamlessness in Tiled Projection-Based Displays", *ACM Trans. on Graphics*, 2005.
- [21] S.K. Nayar, H. Peri, M.D. Grossberg, and P.N. Belhumeur, "A Projection System with Radiometric Compensation for Screen Imperfections", *Proc. of ICCV Workshop on Projector-Camera Systems*, 2003.
- [22] T. Okazaki, T. Okatani, and K. Deguchi, "Shape Reconstruction by Combination of Structured-Light Projection and Photometric Stereo Using a Projector-Camera System", *Pacific Rim Symp. on Advances in Image and Video Technology*, 410-422, 2009.
- [23] J. O'Rourke, "Art Gallery Theorems and Algorithms", *Oxford University Press*, 1987.
- [24] G. Pingali, C. Pinhanez, A. Levas, R. Kjeldsen, M. Podlaseck, H. Chen, and N. Sukaviriya, "Steerable Interfaces for Pervasive Computing Spaces", *IEEE International Conf. on Pervasive Computing and Communications*, 2003.
- [25] R. Raskar, G. Welch, K.L. Low, and D. Bandyopadhyay, "Shader Lamps: Animating Real Objects With Image-Based Illumination", *Proc. of Eurographics Workshop on Rendering Techniques*, 89-102, 2001.
- [26] B. Sajadi, M. Lazarov, A. Majumder, and M. Gopi, "Color Seamlessness in Multi-Projector Displays Using Constrained Gamut Morphing", *IEEE Trans. On Visualization and Computer Graphics*, 2009.
- [27] B. Sajadi and A. Majumder, "Auto-Calibration of Cylindrical Multi-Projector Systems", *IEEE Virtual Reality*, 2010.
- [28] B. Sajadi and A. Majumder, "Markerless View-Independent Registration of Multiple Distorted Projectors on Vertically Extruded Surface Using a Single Uncalibrated Camera", *IEEE Trans. on Visualization and Computer Graphics*, 2009.
- [29] W. Stuerzlinger, "Imaging all Visible Surfaces", *Proc. of Graphics Interface*, 115-122, 1999.
- [30] W. Scott, G. Roth, and J.F. Rivest, "View planning for automated three-dimensional object reconstruction and inspection", *ACM Computing Surveys*, 35:1, 64-96, 2003.
- [31] D. Scharstein and R. Szeliski, "High-Accuracy Stereo Depth Maps Using Structured Light", *IEEE Conf. on Computer Vision and Pattern Recognition*, 195-202, 2003.
- [32] G. Wetzstein and O. Bimber, "Radiometric Compensation Through Inverse Light Transport", *Proc. of Pacific Graphics*, 391-399, 2007.
- [33] P. Wonka, M. Wimmer, K. Zhou, S. Maierhofer, G. Hesina, and A. Reshetov, "Guided Visibility Sampling", *ACM Trans. on Graphics (SIGGRAPH)*, 25:3, 494-502, 2006.
- [34] M. Zwicker and J. Van Baar, "EWA Splatting", *IEEE Trans. on Visualization and Computer Graphics*, 8:3, 223-238, 2002.