

Content and Surface Aware Projection

Long Mai*

Hoang Le†

Feng Liu‡

Portland State University

ABSTRACT

Image projection is important for many applications in entertainment industry, augmented reality, and computer graphics. However, perceived distortion is often introduced by projection, which is a common problem of a projector system. Compensating such distortion for projection on non-trivial surfaces is often very challenging. In this paper, we propose a novel method to pre-warp the image such that it appears as distortion-free as possible on the surface after projection. Our method estimates a desired optimal warping function via an optimization framework. Specifically, we design an objective energy function that models the perceived distortion in projection results. By taking into account both the geometry of the surface and the image content, our method can produce more visually plausible projection results compared with traditional projector systems. We demonstrate the effectiveness of our method with projection results on a wide variety of images and surface geometries.

Index Terms: I.3.8 [Computing Methodologies]: Computer Graphics— Applications;

1 INTRODUCTION

Projectors provide a natural method for displaying visual content on 3D surfaces and have been employed in applications in entertainment industry, augmented reality, and computer graphics. For example, it has been long realized in entertainment industry that extending the field of view of the displayed visual content can greatly enhance the viewing experience. Due to their inherent constraints in physical shape and size, traditional electronic displays often cannot offer such an extreme field of view. Even large-screen TV or large display walls are often limited in creating immersive viewing experience. Many displaying systems have been developed to project visual content onto large, wide field-of-view 3D surfaces, such as the dome theaters [21], flight simulation systems [10], and the Disney’s entertainment theme park projection shows [15]. The ability to project image content to large-size surfaces with any shape makes projectors popular equipment for those applications.

Projectors have also been used in Augmented Reality (AR) applications [2, 8, 19, 20, 28, 29]. Most AR systems augment physical environment with virtual visual content. The ability to flexibly display image content onto 3D environmental surfaces makes projectors good components of such systems. Recently, handheld projector systems have become widely available. For example, many ubiquitous devices such as smart phones or consumer cameras have already embedded pico-projectors. We can therefore expect the use of projector-based systems to display image content to grow even more significantly. In computer graphics, projective texture mapping has been developed as a popular method for automatic texture generation [4]. In projective texturing, the effect of the slide projector is simulated such that 2D texture image is mapped onto geometry in

*e-mail: mtlong@cs.pdx.edu

†e-mail: hoanl@cs.pdx.edu

‡e-mail: fliu@cs.pdx.edu

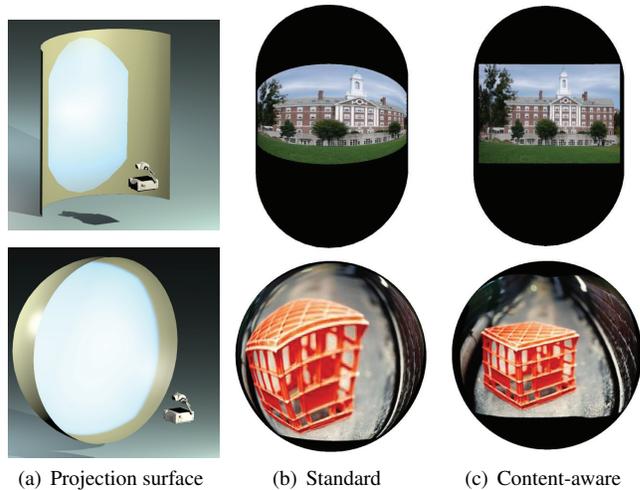


Figure 1: Standard projection vs. content-aware projection. Standard projection induces distortion when projecting images onto non-planar 3D surfaces while content-aware projection can produce a distortion-free result (on the top) or reduce the visually noticeable projection distortion significantly (at the bottom).

graphically generated environments using a projection model. Projective texture mapping has been used in many applications, since it avoids the needs for assigning fixed texture coordinates and provides a good method of representing synthetic images or photographs in image-based rendering [9].

Existing projector systems, however, have a fundamental limitation. If the projector system, including both the projector and projection surface, is not well configured, the projected visual content will appear distorted. That is, the appearance of the projected visual content strictly depends on the projection surface as well as the geometric relationship between the surface and the projector. In fact, for all except a few special configurations, the projected image will be distorted. A well-known problem is the keystone distortion. It appears when the optical axis of a projector is not perpendicular to a planar surface. While automatic keystone correction has received a lot of research effort (e.g. [11, 17, 25]), little effort has been devoted to approaching the more general problem of correcting the projection distortion on a non-planar 3D surface.

In this paper, we present a method that removes or reduces perceived distortion for any viewpoints when projecting an image onto an arbitrary 3D surface. Our solution is to pre-warp the image in such a way that it is distortion-free or least distorted on an arbitrary 3D surface after projection independent from viewpoints. Specifically, we design content-aware distortion metrics that measure how distorted the projection result appears given a warping function. We then estimate a warping function that induces a minimal amount of projection distortion. We experiment with our method by both using computer simulation and developing a prototype content- and surface-aware projector system. Our experiments show that our method can project images onto a variety of 3D surfaces with minimal visual distortion.

2 RELATED WORK

Minimizing distortion from projecting an image onto a 3D surface has been well studied in the fields of computer graphics, computer vision, and multimedia. Most existing methods work on a specific case of projection distortion, keystone distortion. Keystone distortion appears when the optical axis of a projector is not perpendicular to the planar surface. Keystone distortion can be eliminated by either re-orientating the projector or pre-warping the input image using a properly estimated homography transformation [11, 12, 17, 24, 25]. These methods, however, cannot work well when the 3D surfaces are not planar as no homography-based global image warping can account for spatially-varying warping induced from arbitrary 3D surfaces.

Raskar *et al.* developed *iLamps*, a geometry-aware projector system that can make use of a conformal mapping method to pre-warp the input image so that the projection distortion onto a non-planar surface can be reduced [18]. Our work differs from *iLamps* in two aspects. First, the method in *iLamps* only enforces the conformality (i.e. angle-preserving) condition during image warping. Although conformality is necessary to prevent distortion, it alone is often insufficient for general surfaces. Our method takes into account both the angle-preserving and scale-preserving conditions and uses a more complete distortion metric. More importantly, *iLamps* only considers the surface geometry and ignores the image content. For complex 3D surface, projection distortion necessarily occurs. Our method that considers image content can better reduce visually noticeable projection distortion.

This work is also related to the methods that aim to change the original image content such that its projection appears correct when viewed from a predefined target viewpoint [1, 2, 19]. This paper works on a different problem. We aim to warp the image such that the appearance of its projection on the surface is least-distorted independent from any viewpoint. Similar to *iLamps* [18], our goal of viewpoint-independent distortion compensation is to project the image such that the distortion at any local region is small when viewed along the normal vector direction at that region and the whole projected image appear as if it were “pasted” onto the surface.

Our method uses spatially-varying warping to pre-warp an input image. Variations of spatially-varying warping methods have been widely used in computer graphics and computer vision problems, such as image and video re-targeting [27, 30], image manipulation [3, 7, 14], image stitching [31], and video stabilization [13]. Our research works on a different problem and designs a new spatially-varying warping method to minimize projection distortion.

3 CONTENT- AND SURFACE-AWARE PROJECTION

The first question we need to discuss is *what makes a good projection result*. The same projection result is perceived differently when it is viewed from different viewpoints. Ideally, the projector system should project an image differently for different viewpoints. When there are multiple viewers, this is difficult to achieve with a standard projector. In this paper, we focus on the scenario where visual content is projected onto a 3D surface with multiple viewers sitting roughly in front of the projection surface. In this case, we seek such a projection result that appears as minorly distorted to all the viewers as possible although it might not be perfectly distortion-free to any viewer. We consider that a good solution is to project the image onto the 3D surface as if we wrap the image onto the 3D surface with as little non-uniform distortion as possible. When such a projection result is viewed by multiple viewers roughly in front of the surface, it will appear only insignificantly distorted to all of them.

3.1 Overview

In order to achieve the above discussed good projection result, our method pre-warps an input image so that it suffers from no or as

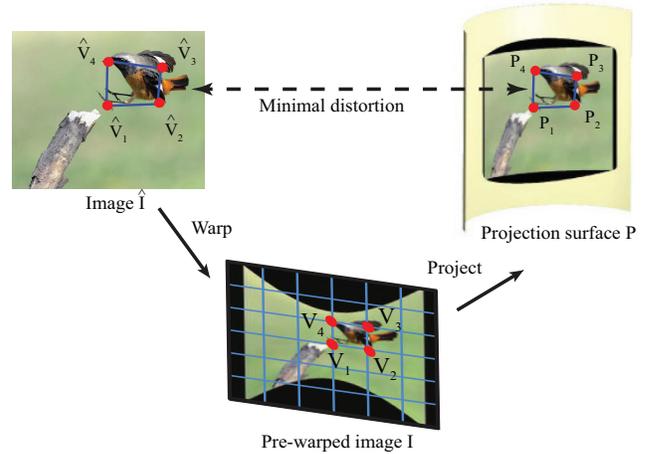


Figure 2: Content- and surface-aware projection framework. Our method pre-warps an input image \hat{I} and projects the pre-warping result I onto 3D surface P . The pre-warping step uses a spatially-varying warping method that warps the input image such that the input image resembles its projection as much as possible.

little as possible distortion when it is projected onto an arbitrary 3D surface. The problem setting of our method is illustrated in Figure 2. Specifically, our method first warps an input image \hat{I} and obtains a pre-warping result I , which is then put onto the projector’s image plane and projected onto 3D surface P .

Our method requires that the projector system is calibrated and the projector surface geometry is known. There have been many methods for projector calibration and surface geometry estimation using a projector system [18]. In our work, we developed a prototype projector-camera system to calibrate the projector and estimate the surface geometry. Our system consists of an Optoma ML750 projector and a Nikon D7000 camera. We employ the projector-camera calibration software package from Moreno *et al.* [16] to calibrate the projector-camera system and use triangulation to recover a sparse proxy of the surface geometry.

The key research question is how to warp input image \hat{I} into I such that the final projection appears with no or minimal distortion. While homography-based global image transformation can handle the common problem of keystone distortion well, it cannot work for more general 3D surfaces than planar surfaces as the projection distortion depends on surface geometry. Thus, we use spatially-varying warping to pre-warp the input image that considers both the image content and 3D surface.

We formulate the image pre-warping problem as a mesh warping problem like previous spatially-varying warping algorithms [3, 13, 26]. Previous methods, however, divide an input image into a uniform grid mesh and optimize an output mesh that meets specific application requirements. We find that this forward-mapping formulation does not suit our problem well. If the mesh unknowns are defined in the pre-warping result, we need to compute the projected mesh positions on the target 3D surface online during mesh optimization. This leads to a time-consuming optimization problem to solve. We instead use a backward-mapping formulation. As shown in Figure 2, we divide the pre-warping result I into a regular grid mesh $V = \{V_i\}$. The goal is to identify for each mesh vertex V_i its optimal corresponding coordinate \hat{V}_i in the original image that the projection of the pre-warping image on 3D surface P resembles the input image as much as possible.

We cast the mesh warping problem as an optimization problem that aims to minimize the projection distortion. We encourage the transformation between the projection of each mesh cell and the corresponding one in the input image to respect a conformal mapping. As we show later on, the conformality constraint alone is often insufficient. We therefore further constrain that the relative

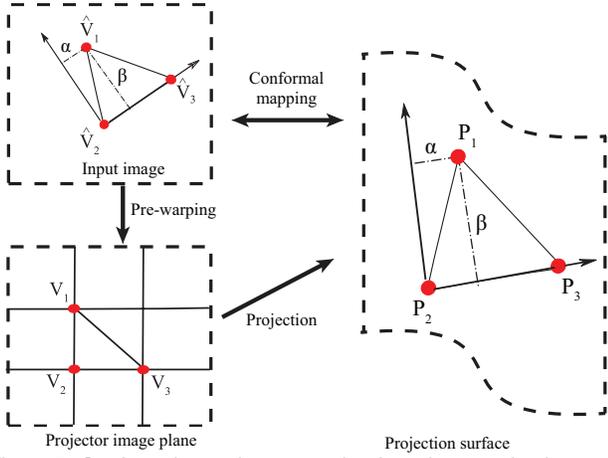


Figure 3: Conformal mapping constraint. In order to make the transformation between the triangle $\hat{V}_1, \hat{V}_2, \hat{V}_3$ and its on-surface projection P_1, P_2, P_3 a conform mapping, the coordinates of \hat{V}_1 in the local coordinate system defined by \hat{V}_2 and \hat{V}_3 must be the same as those of P_1 in the local coordinate system defined by P_2 and P_3 .

size changes among mesh cells should be as similar as possible.

We define the objective energy function as follows

$$E(\hat{V}) = E_a(\hat{V}) + \lambda E_r(\hat{V}), \quad (1)$$

where $\hat{V} = \{\hat{V}_i\}$ is the output mesh in the input image domain (note, we use backward-mapping for the pre-warping step). E_a and E_r represent the energy terms encoding the distortion metrics according to the conformality constraint and relative size change constraint, respectively. λ is a parameter. We discuss the definitions of these energy terms in the following subsections and how to solve the above energy minimization problem in Section 4. Once we obtain the output mesh, we use standard texture mapping to render the pre-warping result [23].

3.2 Conformal Mapping

We define the conformal mapping energy term to preserve the shape of local regions in the original image during projection. Specifically, we divide each grid cell in V into triangles, and then use the method from [7] to enforce the conformality constraint for each of those local triangles.

For example, consider the triangle $\triangle V_1 V_2 V_3$ in the grid mesh of the pre-warping image. We use $\triangle P_1 P_2 P_3$ to denote its projection on the surface (see Figure 3). We note that P_1 can be represented in terms of P_2 and P_3 as follows:

$$P_1 = P_2 + \alpha(P_3 - P_2) + \beta R_{90}(P_3 - P_2), \quad (2)$$

where R_{90} denotes the 90-degree rotation matrix around the normal vector of the plane formed by P_1, P_2 and P_3 , α and β are the local coordinates of P_1 within the coordinate system defined by vector $\overrightarrow{P_2 P_3}$.

In order to preserve the shape of the triangle $\triangle \hat{V}_1 \hat{V}_2 \hat{V}_3$ during projection, the transformation from it to the triangle $\triangle P_1 P_2 P_3$ should be a conformal mapping (i.e. the angle of the triangle should be preserved after the transformation). According to Igarashi *et al.* [7], such condition implies that \hat{V}_1 be represented in term of \hat{V}_2 and \hat{V}_3 as

$$\hat{V}_1 = \hat{V}_2 + \alpha(\hat{V}_3 - \hat{V}_2) + \beta R_{90}(\hat{V}_3 - \hat{V}_2) \quad (3)$$

where α and β are the known coordinates computed from Equation (2). We accordingly define the local angle-preserving-based distortion measurement at the vertex V_1 as follows

$$E_a(\hat{V}_1) = w_a(\hat{V}_1) \|\hat{V}_1 - \hat{V}_2 - \alpha(\hat{V}_3 - \hat{V}_2) - \beta R_{90}(\hat{V}_3 - \hat{V}_2)\|^2. \quad (4)$$

The full energy term $E_a(\hat{V})$ can then be defined as the summation of Equation (4) over all eight triangles for all mesh vertices V_i . As surveyed in Floater and Hormann [5], other alternative formulations of conformal mapping are available. We used the above one as it only involves quadratic terms and is widely used in computer graphics [7].

As the projection distortion necessarily occurs for complex 3D surface, standard conformal mapping tends to distribute the distortion over the whole surface. Our method distributes distortion less to salient regions than those less salient ones. As a result, our method encourages small distortion in important regions with the trade off of having large distortion in less important region such as sky or grass.

We use a weighting factor w_a to modulate the above energy term. To compute w_a , our method first automatically estimates a saliency map S_I from the input image using the graph-based saliency detection method [6]. Our method also allows the user to manually identify the importance map by roughly annotating the most important regions in the image, resulting in a binary mask M_I . We combine the saliency map S_I and the user mask M_I pixel-by-pixel using the max operator to obtain the final content map C . The weighting factor w_a can then be computed from the content map C as

$$w_a(\hat{V}_i) = \delta(C(\hat{V}_i) \geq T) \quad (5)$$

where δ is the indicator function which outputs 1 when its argument is true and 0 otherwise. T is a threshold value. We experimentally determine the value of $T=0.67$ to leverage both the saliency map and the user annotation map when they disagree. Using too low the value of T leads to the noisy content map where unimportant textured regions, such as grass and water, may get high weights. On the other hand, using too large the value of T may cause the user annotation map to overly dominate, which may miss important regions that were not masked by the user. We found that for the saliency detection method used in this paper, the value of 0.67 is sufficient to remove the noise the in saliency maps.

3.3 Relative-size Preserving

While the above conformality constraint can preserve the shape of each local mesh cell well, it alone is often not sufficient to minimize the projection distortion. Although the shapes of local image regions are preserved through conformal mapping, the change in relative size across regions leads to clear distortion in the projection result. To further reduce the distortion, the relative sizes between different regions should also be preserved during projection. We define the energy term E_r below to encode such a constraint.

We use A_i to denote the four grids around the grid point i and $\hat{V}(A_i)$ and $P(A_i)$ to denote the corresponding region of A_i in the input image and its projection on the surface, respectively. $s(\hat{V}_i)$ is the scaling factor of the transformation from $\hat{V}(A_i)$ to $P(A_i)$. The relative-size-preserving energy term is defined as follows:

$$E_r(\hat{V}) = \sum_{i,j} (w_a(\hat{V}_i) + w_a(\hat{V}_j)) \|s(\hat{V}_i) - s(\hat{V}_j)\|^2, \quad (6)$$

where $w_a(\hat{V}_i)$ and $w_a(\hat{V}_j)$ are saliency values at V_i and V_j as defined in Equation (5).

Figure 4 illustrates the effects of different projection methods with a synthetic example where the input image is covered by a regular grid. An artificial content map is created to cover a portion of the image. The results indicate that combining the conformal mapping constraint with the size-preserving constraint can successfully keep the grid cells in the masked region square and of similar sizes by distorting the cells in other regions. Using the conformal mapping constraint alone, on the other hand, cannot preserve the relative size among the grid cells in the masked region and make the region as a whole appear distorted.

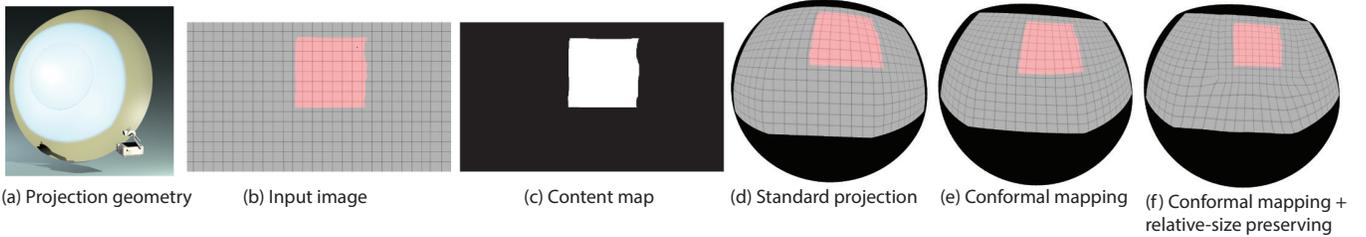


Figure 4: Synthetic example. The input image is covered by a regular grid. The content map is artificially created to cover a portion of the image. Combining the conformal mapping constraint with the size-preserving constraint can successfully keep the grid cells in the masked region square and have similar sizes to each other. Using conformal mapping alone cannot preserve the relative size among the grid cells and make the whole masked region appear distorted.

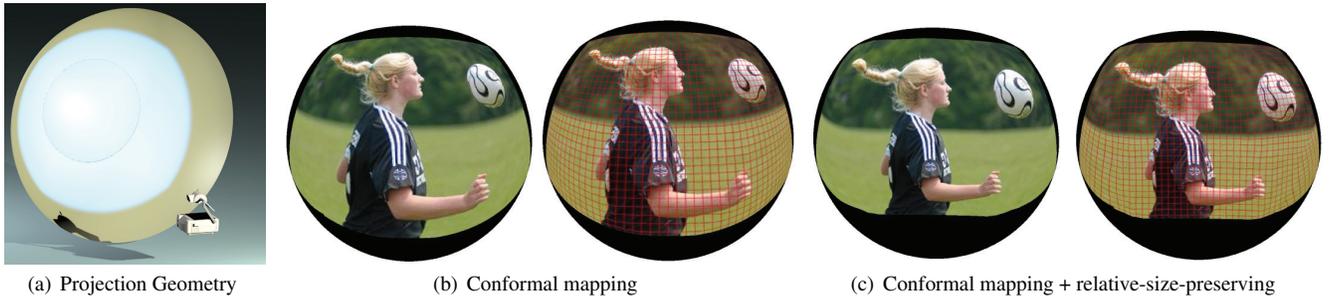


Figure 5: Relative-size-preserving constraint. When images are projected onto a complex 3D surface like a bowl-shaped surface, conformal mapping alone often cannot effectively reduce the projection distortion, as the sizes of different image regions are changed inconsistently. As shown in (b), the salient image content, the girl, is unnaturally distorted even though each cell in the salient image region respects a conformal mapping well. Combining the relative-size-preserving constraint can better reduce the projection distortion. As shown in (c), the mesh cells in the salient image regions are changed consistently and the distortion is distributed more in less salient regions.

Figure 5 further illustrates the effect of incorporating the relative-size-preserving constraint with a real example. As shown in Figure 5(b), the projected image of the girl appears distorted as the sizes of her head and her body are changed differently. Combining the relative-size-preserving constraint can better reduce the projection distortion. As shown in Figure 5(c), the mesh cells in the salient image regions are changed consistently and the distortion is distributed more in less salient regions.

Figure 6 shows an example of projection results generated by different projection methods, viewed from different viewpoints. By incorporating both conformal mapping and relative-size-preserving constraints in a content-aware manner, our method can encourage as minimal distortion as possible in important regions and at the same time preserve the object size in these regions consistently, which gives more natural projection results compared to both the standard projection and the projection with only the conformal mapping constraint.

4 OPTIMIZATION

We note that the values of w_a are determined through indexing into the saliency map using the variables \hat{V} being optimized. For that reason, the overall energy function $E(\hat{V})$ in Equation (1) is highly non-linear. On the other hand, if the values for w_a are fixed, $E(\hat{V})$ can be optimized much easier. We therefore minimize $E(\hat{V})$ using an alternation optimization technique. Specifically, we first fix w_a values and optimize the reduced problem of $E(\hat{V})$, and then update w_a values using the optimization result. We iterate this process until the solution converges. Our algorithm is outlined in Algorithm 1.

4.1 Initialization

To start the alternation optimization process, our method estimates the initial solution by minimizing

$$E_0(\hat{V}) \equiv E'(\hat{V}), s.t. \hat{V}_0 = (0, 0), \hat{V}_{|V|} = (1, \gamma), \quad (7)$$

where $E'(\hat{V})$ is equivalent to $E(\hat{V})$ with all $w_a(\hat{V}_i)$ replaced by $(w_a)_i = 1$. We assume that the input image space has been normalized to the range $[0, 1] \times [0, \gamma]$, where γ is the aspect ratio of the input image. The boundary constraints are used to fix the corresponding coordinates of the two corner vertices at the corners of the input image. Otherwise, the energy function above is under-determined. We use the same method described in the next subsection to solve Equation (9) to solve this minimization problem.

The solution \hat{V}' obtained from solving $\hat{V}' = \operatorname{argmin}_{\hat{V}} E_0(\hat{V})$ may not provide a good initialization as it may cause the loss of image content after projection. This can happen when some boundary vertex V_i is mapped into the interior part of the image (i.e. $\hat{V}'_i \in [0, 1] \times [0, \gamma]$). To further improve this initial solution, we transform the coordinates of every \hat{V}' using a similarity transformation such that the rectangle $R \equiv [0, 1] \times [0, \gamma]$ becomes the maximal inscribed rectangle of the bounded polygon formed by the boundary vertices. That can be equivalently achieved by searching for the maximal inscribed rectangle R' of the boundary polygon formed by \hat{V}' with the same orientation and aspect ratio as the rectangle R and then transforming every \hat{V}'_i from the coordinate system of R to that of R' . The transformed coordinates \hat{V}'_i can be used to initialize the alternation optimization process described below.

4.2 Alternation Process

Given the solution $\hat{V}^{(n)}$ at the n^{th} iteration, we first update the saliency weighting factors w_a as

$$(w_a)_i = w_a(\hat{V}_i^{(n)}) = \delta(C(\hat{V}_i^{(n)}) \geq T). \quad (8)$$

We then obtain the solution at the $(n+1)^{\text{th}}$ iteration by minimizing

$$E_{n+1}(\hat{V}) \equiv E'(\hat{V}) + \sum_i (1 - C(\hat{V}_i)) \|\hat{V}_i - \hat{V}_i^{(n)}\|^2, \quad (9)$$

$$s.t. \quad \hat{V}_i = \hat{V}_i^{(0)}, \forall V_i \in B(V)$$

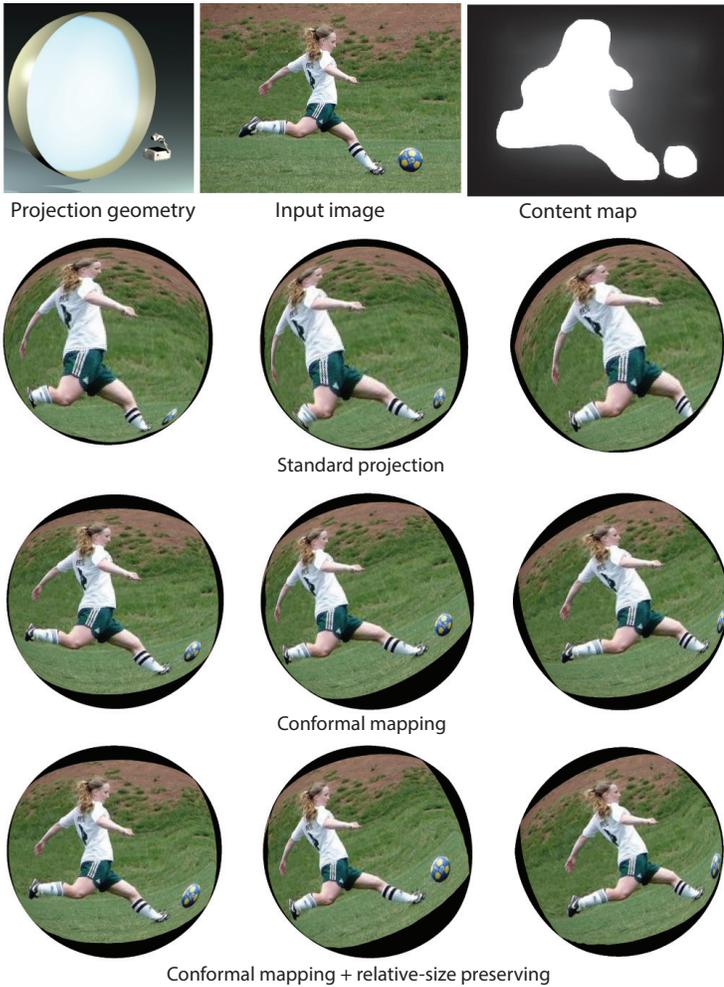


Figure 6: Projection results generated by different projection methods, viewed from different viewpoints. Our method incorporates both conformal mapping and relative size-preserving constraints in a content-aware manner to encourage small distortion in local regions and at the same time preserve global size consistency.

The second term in the above equation is a regularization term which is designed to smooth the transition between solutions of consecutive iterations. It prevents the new solution from deviating faraway from the previous one for the less salient regions. $B(V)$ denotes the set of boundary vertices in the grid mesh. This hard boundary constraint is added to fix the boundary to that from the initial solution $\hat{V}_i^{(0)}$ and it prevents the boundary mesh vertices from mapping onto the interior area of the original image and thus avoids content loss after projection.

We now elaborate how to minimize the energy functions defined in Equation (7) and Equation (9). We note that the conformality constraint term E_a in these energy functions is quadratic and can be easily minimized. But the relative-sizing-preserving term E_r in these energy functions involves the scaling factor $s(\hat{V}_i)$ and makes these energy functions non-quadratic and difficult to optimize. To address this problem, we create an auxiliary variable s_i for each $s(\hat{V}_i)$ and introduce another term to the optimization function to relate each s_i to its corresponding \hat{V}_i

$$E_r(\hat{V}) = \sum_{i,j} (w_a(\hat{V}_i) + w_a(\hat{V}_j)) \|s_i - s_j\|^2 + \sum_{\hat{V}_i} r(\hat{V}_i, s_i) \quad (10)$$

$$r(\hat{V}_i, s_i) = \sum_{j \in N_i} \|\hat{V}_i - \hat{V}_j - s_i l(P_i, P_j) \vec{v}_{ij}^{\rightarrow}\|^2 \quad (11)$$

Algorithm 1: Alternation optimization. Specifically, we first fix the saliency weights $(w_a)_i$ and minimize the energy function E in Equation (1) and then update $(w_a)_i$ using the minimizing result. We repeat this procedure until the optimization process converges.

input : $P = \{P_i\}$
output : \hat{V}^* that minimizes $E(\hat{V})$

- 1 $(w_a)_i \leftarrow 1, \forall i;$
- 2 $\varepsilon \leftarrow \infty$ // residual value used to evaluate convergence condition
- 3 $iter \leftarrow 0;$
- // Initialization
- 4 Compute the initial solution $\hat{V}^{(0)}$ as described in Section 4.1
- // Alternation process
- 5 **while** $\varepsilon > tolVal$ and $n < maxIter$ **do**
- 6 $(w_a)_i \leftarrow \delta(S(\hat{V}_i^{(n)})) \geq T, \forall i;$
- 7 $\hat{V}^{(n+1)} \leftarrow \operatorname{argmin}_{\hat{V}} E_{n+1}(\hat{V});$
- 8 $n \leftarrow n + 1;$
- 9 $\varepsilon \leftarrow \frac{1}{|\hat{V}|} \sum_i \|\hat{V}^{(n)} - \hat{V}^{(n-1)}\|;$
- end**
- 10 $\hat{V}^* \leftarrow \hat{V}^{(n)};$

where $l(P_i, P_j)$ is the length of the vector $\vec{V}_i \vec{V}_j$ after projection, and N_i denotes the set of vertices adjacent to V_i in the grid mesh. $\vec{v}_{ij}^{\rightarrow} = \frac{\hat{V}_i - \hat{V}_j}{\|\hat{V}_i - \hat{V}_j\|}$ denotes the unit vector indicating the orientation from \hat{V}_i to \hat{V}_j . Intuitively, $r(\hat{V}_i, s_i)$ enforces the local region around \hat{V}_i to have the same scale factor s_i .

Directly using the $\vec{v}_{ij}^{\rightarrow}$ direction vector as defined above introduces the non-linear factor due to the computation of the vector norm. We follow the two-step approach introduced by [7] to address this problem. In the first step, we solve the optimization problem given by Equation (1) with λ_2 set to 0 (equivalent to ignoring the E_r energy term) using a standard sparse linear solver. The solution \hat{V}' obtained from that will be used to compute $\vec{v}_{ij}^{\rightarrow}$ as

$$\vec{v}_{ij}^{\rightarrow} = \frac{\hat{V}'_i - \hat{V}'_j}{\|\hat{V}'_i - \hat{V}'_j\|}. \quad (12)$$

Once $\vec{v}_{ij}^{\rightarrow}$ is fixed, the energy terms resulting from $r(\hat{V}_i, s_i)$ will only contribute linear least-square terms to the whole energy function. The resulting optimization problem thus becomes a linear least-square problem and can then be efficiently solved in the second step using a sparse linear solver. Note that although the variable s_i 's are also included in the optimization problem, they are just the auxiliary variables. After the optimization process, we only make use of the value of \hat{V}'_i 's in determining the warping function.

The alternating process is repeated until convergence or until it reaches the maximum allowed number of iterations. Figure 7 shows an example where the projection distortion keeps decreasing during iteration. In our experiments, we observed that typically less than 10 iterations are needed to produce visually pleasing results.

5 EXPERIMENTS

We experimented with our content- and surface-aware projection method on a range of 3D projection surfaces and compared our method to the standard projection method. We also discuss components and variations of our method. Please refer to the video demo



Figure 7: Iteration results. The projection results are improved over iterations.

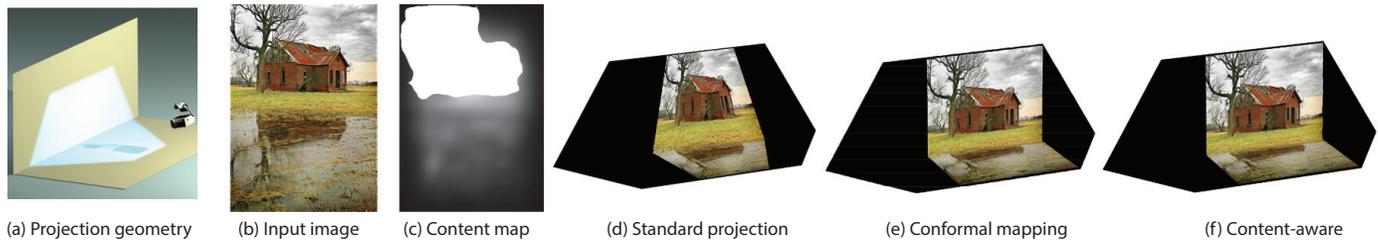


Figure 8: Developable surface. For developable surfaces, conformal-mapping-based projection and content-aware projection can both produce distortion-free results while standard projection introduces distortion. The content map is obtained from the combination of automatically generated saliency map and user provided annotation mask as described in Section 3.2.

in our project website to better examine our method, our prototype system, and results¹.

5.1 Developable Surfaces

We first examine our method on a special category of 3D surfaces, namely developable surfaces. A developable surface is mathematically defined as a surface with zero Gaussian curvature. Informally speaking, a surface is called developable if it can be “flattened” into a plane without stretching or compressing. Figure 8 shows an example of projecting an image onto the intersection corner of two walls. On this surface, we find that the standard projection method distorts the images, as shown in Figure 8 (d). We note that the content map used in our methods are obtained from the combination of automatically generated saliency map and user provided annotation mask, as described in Section 3.2.

It is interesting to observe that both conformal mapping and content-aware projection produce distortion-free results, as shown in Figure 8 (e) and (f). That is because a distortion-free conformal mapping exists for a developable surface. In fact, that mapping is not only conformal (angle-preserving) but also authalic (i.e. scales and distances are also preserved) [22]. For these types of surfaces, there is no conflict between the angle-preserving constraint and the relative-size preserving constraint (the energy terms E_a and E_r respectively). As a result, for the case of developable surfaces, our content-aware projection method automatically generates the same distortion-free projection result as the conformal-mapping-based method.

5.2 Non-developable Surfaces

For surfaces that are not developable (e.g. the sphere, the bowl-shaped surface), it was known that there exists no mapping which is both conformal and authalic [22]. In other words, in order to obtain a mapping with little angle distortion, it is necessary to trade off the relative-size change constraint and thus the projection distortion is inevitable. Figure 9 shows such an example where an image is projected onto a surface formed by a half cylinder and a floor. We can see that while conformal mapping alone introduces distortion particularly in the light house region, combining the relative-size

change constraint cannot avoid distortion either. Our method considers image content and distributes more distortion to those less salient regions, thus minimizing the visually noticeable distortion, as shown in Figure 9 (f).

Figure 10 shows more non-developable surfaces. The first row shows an example where a leopard image is projected onto the corner of a cylindrical wall and a floor. Enforcing the conformal-mapping constraint leads to the inconsistent size change among different parts of the leopard after projection. For example, the leopard’s head becomes unnaturally smaller than its body, which leads to the noticeable visual distortion in the projection result (Figure 10 (e)). The projection distortion is significantly reduced using our content-aware projection method, as shown in Figure 10 (f).

The second row in Figure 10 shows another non-developable surface. In this example, an image is projected onto a spherical surface. We observe that conformal-mapping changes the sizes of two persons inconsistently. Our content-aware projection method resolves the problem and makes the projection result appear as if a poster were “pasted” onto the surface.

5.3 A Real Projection System

To further evaluate our projection method, we developed a projector-camera system which consists of an Optoma ML750 projector and a Nikon D7000 camera. We use the projector-camera calibration software package from Moreno *et al.* [16] to calibrate the projector-camera system and use triangulation to compute the 3D coordinates for each grid point in the projector image plane and use them as the input for our image pre-warping method described previously.

Figure 11 shows our projection results on real 3D surfaces. The figure also shows the standard projection results (i.e. projection of the original input image without pre-warping) for comparison. From the figure, we can see that our projection method can significantly reduce visual distortion when projecting image content onto those non-trivial surfaces. Note that the projected image content appears as if it was “pasted” onto the surface.

5.4 Discussions

Our method currently does not include the constraints for the projection of image boundary. The boundary in the projection result

¹<http://graphics.cs.pdx.edu/project/adaproj>

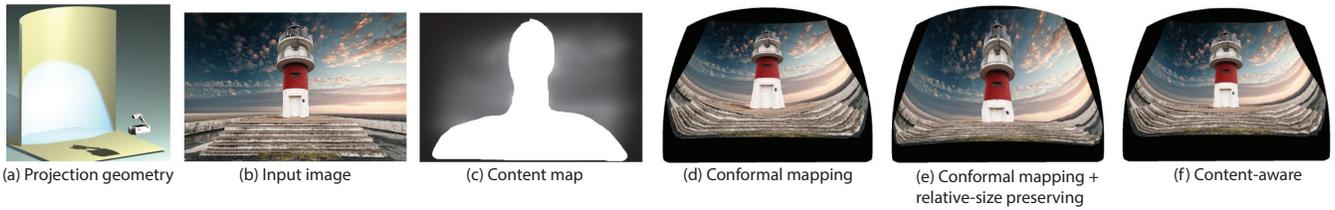


Figure 9: Content-aware projection. For a non-developable surface, conformal mapping, even with the relative-size-preserving constraint, cannot avoid the projection distortion. Content-aware projection can better preserve salient content by distributing more distortion to less salient ones. As shown in (f), the light house is less distorted using content-aware projection than other methods.



Figure 10: Examples of projection on non-developable surfaces. These examples show that our content-aware method can preserve the shape of the important image content well by distributing more projection distortion to less semantically important regions like sky, grassland, or blurry background.

sometimes may appear unattractive. Meanwhile whether the appearance of the projection boundary is pleasant or not is often a question of art and application-dependent, especially when an image is projected onto an arbitrary 3D surface. Our method currently leaves these irregular boundaries. If needed, we can crop off the irregular boundaries or incorporate additional boundary constraints into our optimization framework.

The major computation cost of our method is in the pre-warping step. The computational cost depends on the surface properties and the mesh size. We report the computational cost for a typical mesh size 20×20 below on a computer with a CORE i5 1.7GHZ CPU and 6GB memory. For developable surfaces, our method only requires the initialization step, about 0.01 seconds, to produce distortion-free projection results. For non-developable surfaces, it typically takes less than 10 iterations, with each iteration taking about 0.3 seconds. This limits our method from applying to video projection. If we do not consider image content, the pre-warping is fixed and thus only needs to be computed once and off-line. This means that the warping is fixed for all the frames. This simplification makes our method applicable for video at the cost of projection quality. Our current method needs to be further sped up to project a video onto a non-developable surface in realtime and in a content-aware way.

6 CONCLUSION

In this paper, we described a content-aware projection method to project visual content onto arbitrary 3D surfaces. Our method employs a content-aware warping method to pre-warp an input image such that the visual projection distortion can be minimized. This pre-warping method considers both the image content and the sur-

face geometry. Our experiments show that our method can produce distortion-free projection results for developable surfaces and significantly reduce visual distortion for non-developable surfaces. In future, we will improve the speed of our method so that our content-aware projection can be used to project streaming visual content, like video, in realtime.

ACKNOWLEDGMENTS

The images in Figures 2, 5, 6, 7, 9, the second row of Figure 1, the second row of Figure 10, and the first and third rows of Figure 11 are used under Creative Commons licenses from Flickr users including Mugley, Zachstern, Caese, Pensiero, JohnFish, Danielgilles, Steven-depolo, and CLF. The image in the second row of Figure 11 is used under a Creative Commons license from the Blender Foundation.

REFERENCES

- [1] O. Bimber, A. Emmerling, and T. Klemmer. Embedded entertainment with smart projectors. *Computer*, 38(1):48–55, 2005.
- [2] O. Bimber and R. Raskar. *Spatial augmented reality - merging real and virtual worlds*. A K Peters, 2005.
- [3] R. Carroll, M. Agrawal, and A. Agarwala. Optimizing content-preserving projections for wide-angle images. *ACM Transactions on Graphics*, 28(3):43:1–43:9, 2009.
- [4] C. Everitt. Projective texture mapping, 1999.
- [5] M. S. Floater and K. Hormann. Surface parameterization: a tutorial and survey. In *Advances in Multiresolution for Geometric Modelling*, pp. 157–186, 2005.
- [6] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In *Advances in Neural Information Processing Systems*, pp. 545–552, 2006.

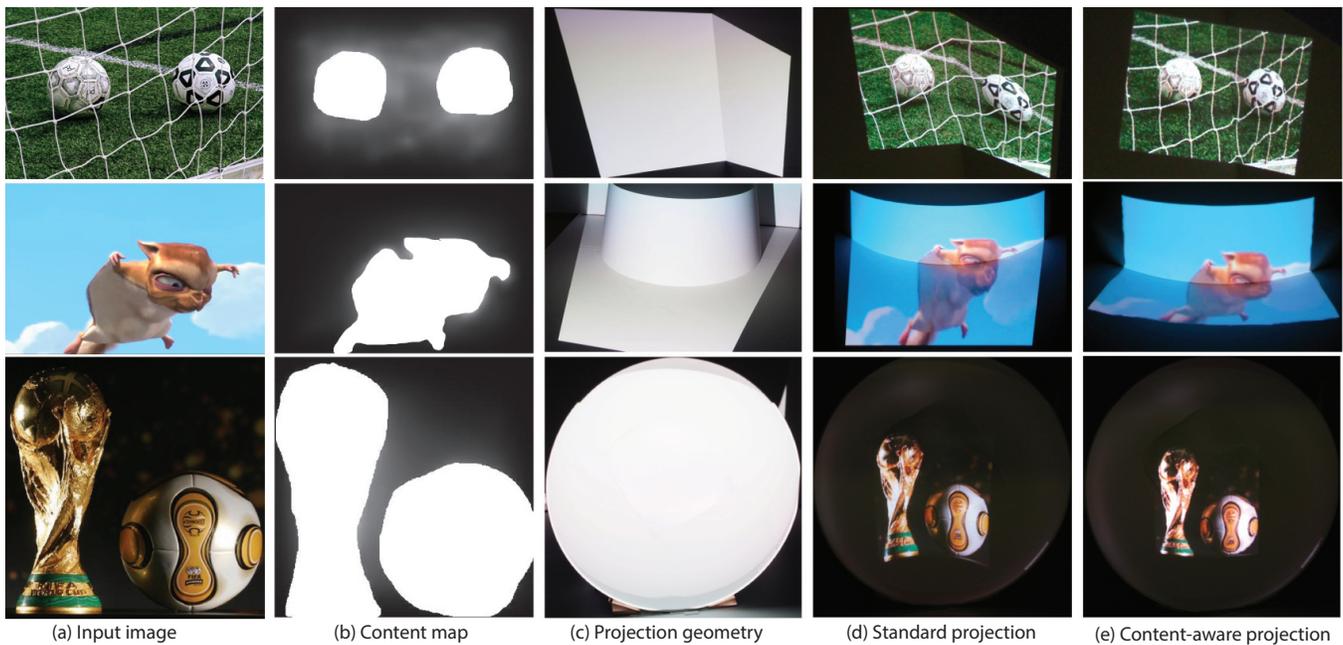


Figure 11: Examples of projection on real 3D surfaces. Our content-aware projection method can significantly reduce visual distortion when projecting image content onto non-trivial surfaces such as a cylindrical surface (first row), an intersection of two planes (second row), an intersection between a plane and a conical surface (third and fourth row), and a spherical surface (fifth row).

- [7] T. Igarashi, T. Moscovich, and J. F. Hughes. As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics*, 24(3):1134–1141, 2005.
- [8] B. R. Jones, H. Benko, E. Ofek, and A. D. Wilson. Illumiroom: peripheral projected illusions for interactive experiences. In *ACM CHI Conference on Human Factors in Computing Systems*, pp. 869–878, 2013.
- [9] D. Kim and J. K. Hahn. Projective texture mapping with full panorama. *Computer Graphics Forum*, 21(3):421–430, 2002.
- [10] M. Lacroix. A hdtv projector for wide field of view flight simulators. In *IMAGE VI Conference*, pp. 493–500, 1992.
- [11] B. Li and I. Sezan. Automatic keystone correction for smart projectors with embedded camera. In *IEEE International Conference on Image Processing*, vol. 4, pp. 2829–2832, 2004.
- [12] Z. Li, K. Wong, Y. Gong, and M. Chang. An effective method for movable projector keystone correction. *IEEE Transactions on Multimedia*, 13(1):155–160, 2011.
- [13] F. Liu, M. Gleicher, H. Jin, and A. Agarwala. Content-preserving warps for 3d video stabilization. *ACM Transactions on Graphics*, 28(3):44:1–44:9, 2009.
- [14] S.-J. Luo, I.-C. Shen, B.-Y. Chen, W.-H. Cheng, and Y.-Y. Chuang. Perspective-aware warping for seamless stereoscopic image cloning. *ACM Trans. Graph.*, 31(6):182:1–182:8, 2012.
- [15] M. Mine, D. Rose, B. Yang, J. van Baar, and A. Grundhofer. Projection-based augmented reality in disney theme parks. *Computer*, 45(7):32–40, 2012.
- [16] D. Moreno and G. Taubin. Simple, accurate, and robust projector-camera calibration. In *Proceedings of the 2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, 3DIMPVT '12, pp. 464–471, 2012.
- [17] R. Raskar and P. A. Beardsley. A self-correcting projector. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 504–508, 2001.
- [18] R. Raskar, J. van Baar, P. A. Beardsley, T. Willwacher, S. Rao, and C. Forlines. ilamps: geometrically aware and self-configuring projectors. *ACM Transactions on Graphics*, 22(3):809–818, 2003.
- [19] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: a unified approach to image-based modeling and spatially immersive displays. In *ACM SIGGRAPH*, pp. 179–188, 1998.
- [20] R. Raskar, G. Welch, K.-L. Low, and D. Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In *Eurographics Workshop on Rendering Techniques*, pp. 89–102, 2001.
- [21] J. Shaw and E. Lantz. Dome theaters: Spheres of influence. In *TILE Proceedings*, pp. 59–65, 1998.
- [22] A. Sheffer, E. Praun, and K. Rose. Mesh parameterization methods and their applications. In *Foundations and Trends in Computer Graphics and Vision*, p. 2006. Now Publishers, 2006.
- [23] P. Shirley and S. Marschner. *Fundamentals of Computer Graphics*. A. K. Peters, Ltd., Natick, MA, USA, 3rd ed., 2009.
- [24] R. Sukthankar, R. G. Stockton, and M. D. Mullin. Automatic keystone correction for camera-assisted presentation interfaces. In *ACM International Conference on Advances in Multimodal Interfaces*, pp. 607–614, 2000.
- [25] R. Sukthankar, R. G. Stockton, and M. D. Mullin. Smarter presentations: Exploiting homography in camera-projector systems. In *IEEE International Conference on Computer Vision*, pp. 247–253, 2001.
- [26] Y.-S. Wang, H.-C. Lin, O. Sorkine, and T.-Y. Lee. Motion-based video retargeting with optimized crop-and-warp. *ACM Transactions on Graphics*, 29(4):90:1–90:9, 2010.
- [27] Y.-S. Wang, C.-L. Tai, O. Sorkine, and T.-Y. Lee. Optimized scale-and-stretch for image resizing. *ACM Transactions on Graphics*, 27(5), 2008.
- [28] K. D. Willis and I. Poupyrev. Motionbeam: designing for movement with handheld projectors. In *ACM CHI Extended Abstracts on Human Factors in Computing Systems*, pp. 3253–3258, 2010.
- [29] K. D. Willis, I. Poupyrev, S. E. Hudson, and M. Mahler. Side-by-side: ad-hoc multi-user interaction with handheld projectors. In *ACM Symposium on User Interface Software and Technology*, pp. 431–440, 2011.
- [30] L. Wolf, M. Guttman, and D. Cohen-or. Non-homogeneous content-driven video-retargeting. In *IEEE International Conference on Computer Vision*, pp. 1–6, 2007.
- [31] J. Zaragoza, T.-J. Chin, M. S. Brown, and D. Suter. As-projective-as-possible image stitching with moving dlt. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2339–2346, 2013.