

Relighting Acquired Models of Outdoor Scenes

Alejandro Troccoli

Peter K. Allen

*Department of Computer Science
Columbia University, New York*

Abstract

In this paper we introduce a relighting algorithm for diffuse outdoor scenes that enables us to create geometrically correct and illumination consistent models from a series of range scans and a set of overlapping photographs that have been taken under different illumination conditions. To perform the relighting we compute a set of mappings from the overlap region of two images. We call these mappings Irradiance Ratio Maps (IRMs). Our algorithm handles cast shadows, being able to relight shadowed regions into non-shadowed regions and vice-versa. We solve these cases by computing four different IRMs, to handle all four combinations of shadowed vs non-shadowed surfaces. To relight the non-overlapping region of an image, we look into the appropriate IRM which we index on surface normal, and apply its value to the corresponding pixels. The result is an illumination consistent set of images.

1. Introduction

Our experience in modeling architectural and archaeological sites [2, 3] has taught us that accurately capturing the appearance of an outdoor scene is a difficult process. Ideally, one would like to take a set of range scans and photographs to produce a correct geometric and photometric model. But in practice, modeling is an incremental process in which range scan and image acquisition are interleaved. The main consequence of this interleaving is that the images that we capture to represent the appearance of the scene are spaced in time, and hence, since our sites are outdoors, the illumination conditions can change from photograph to photograph. Thus, these images no longer represent a consistent appearance of the scene. If they are used without further processing to render a textured model, the result will show noticeable color discontinuities. In this paper we address this problem and show how to relight a set of photographs of an outdoor scene taken under different illumination conditions to generate an illumination consistent appearance of the model.

The relighting algorithm uses the region of overlap be-

tween two photographs to compute a set of mappings that we will use to relight the non-overlapping region. By exploiting the fact that most of the surfaces on our scenes are diffuse and that illumination is distant, these mappings, which we will refer to as Irradiance Ratio Maps (IRM), are computed per surface normal. Shadows cast by the sun will be present, and our proposed algorithm will handle them. To relight in the presence of shadows we will need to compute four IRMs, one that will map non-shadowed regions to non-shadowed regions, one that will map shadowed regions to shadowed regions, and two more, for the non-shadowed versus shadowed cases. Once the IRMs are computed, each pixel in the source image is multiplied by its corresponding irradiance ratio to obtain the relighted color.

Relighting of outdoor scenes is a difficult problem and little previous work has been done in the area. These scenes consist of complex geometry and contain numerous occluding surfaces that make the problem even harder. We present relighting results for a scanned building on the Columbia University campus, and highlight the key aspects of our algorithm along with areas for future improvement.

2. Background and motivation

Our long-term goal is to build accurate geometric and photometric models of large sites that can be used for visualization in augmented reality applications and also serve as teaching tools. In archaeology, for instance, an augmented reality application can assist researchers in post-excavation analysis. During the past years we have been developing a set of tools to automate the modeling pipeline, with emphasis in point cloud registration and image registration. We introduced in [16] a set of algorithms for planar segmentation of point clouds, 3D line extraction and texture registration. The extracted 3D lines can be used for point cloud registration, as presented in [3]. For scenes that do not contain a large number of 3D line segments and planar regions, we have developed a method that uses the shadows cast by the sun for image registration [19]. We have successfully applied these ideas to a real archaeological excavation conducted by a team from the Stanford Archaeology Center. The complete 3D modeling and visualization pipeline is de-

scribed in [2].

3. Previous work

Modeling geometry and appearance has been an area of active research during the past decade. Researchers have used photogrammetry and range sensors to model geometry and photographs to capture appearance. In addition, different techniques and algorithms have been developed to combine these photographs into a suitable representation for the appearance of the scene or object.

Appearance attributes have been represented in different ways. Early work used an ad-hoc weighted average of an input set of images [8, 14, 7]; later, more sophisticated illumination invariant representations were sought [9, 22, 15, 6]. In these works the images are passed to an algorithm that computes the reflectance properties of the scene. The computed surface reflectance constitute an illumination invariant representation for the appearance attributes of the scene. When considering different representation for appearance, there is no such thing as a 'correct' representation. Instead, the right representation is the one that suits the needs of the underlying application. Illumination invariant representations of appearance are good for graphics application where an object or scene needs to be rendered under different illumination conditions. In our visualization applications we do not require the illumination to change; we are only interested in displaying our models in a static illumination environment, but we require the appearance attributes to be consistent with that illumination. We call such a representation an illumination consistent texture.

3.1. Illumination invariant textures

An illumination invariant texture is intrinsic to the surface of the scene. This is a desirable appearance representation for scenes that need to be rendered under any illumination condition. To obtain an illumination invariant texture from images, one must factor the illumination from the surface reflectance function that generated the input image set.

The bidirectional reflectance distribution function (BRDF) serves as an illumination invariant appearance representation. It describes how a surface reflects the incident light. It is a five variable function $f_r(\theta_i, \phi_i, \theta_o, \phi_o, \lambda)$ that defines the ratio of incident to reflected radiance at wavelength λ , for a pair of incident and viewing directions (θ_i, ϕ_i) and (θ_o, ϕ_o) respectively. The recovery of the BRDF from a set of images requires some knowledge of the image formation process. When a camera images a surface point \mathbf{x} on an object, the pixel intensity I corresponding to \mathbf{x} is given by:

$$I(\mathbf{x}) = s(\lambda) \int_{\Omega_i} L(\mathbf{x}, \theta_i, \phi_i, \lambda) f_r(\mathbf{x}, \theta_i, \phi_i, \theta_o, \phi_o, \lambda) \cos \theta_i d\omega_i \quad (1)$$

where $s(\lambda)$ is the camera response function, L the incident radiance, i.e. illumination, f_r is the BRDF and (θ_o, ϕ_o) is the direction of the viewer from \mathbf{x} in local coordinates.

The high dimensionality of the BRDF has encouraged researchers to seek efficient and more compact representations that make the recovery process easier. By ignoring the dependence on wavelength, the BRDF can be reduced to a 4 variable function. In addition, if the BRDF is isotropic, a further reduction can be made to make the BRDF a 3D function. However, even for a 3D function, the number of measurements required to completely acquire the BRDF of an object is large. Further simplifications can be done if one can assume that the reflectance on a surface point follows a certain model that depends on a small number of parameters. The simplest of such models is the Lambertian reflectance model, which states that light is scattered equally in all outgoing directions and hence the BRDF is a constant. Other more realistic and physically possible models include Ward's model for anisotropic reflection [21], Laforune's model for modeling specularities at different reflection angles [9], and Oren and Nayar's model for rough diffuse objects [12], among others.

The problem of recovering the BRDF from images has received much attention. It requires inverting the image formation process described by equation (1) and solving for f_r . The input to this process is a set of images with known camera parameters and position, the geometry of the scene and a measurement of the incident illumination. The output is the 4D BRDF or one of its parametric versions. A number of researchers have described methods for fitting reflection models to measured sample data under laboratory and controlled illumination conditions (usually point light sources). [9, 10, 22, 15, 6].

In an outdoor setting illuminated by the sun and sky, the case that interests us most, solving this problem involves a more complex process. First, it is necessary to capture the complete incident illumination. This is not trivial for the sun and sky, which together have a dynamic range far greater than that of an indoor setting. Love [11] attempts a simplified version of this problem. Keeping the camera fixed, he acquires a set of images of a flat sample at different times of the day. He uses a sky and sun model to compute the incident illumination and solves for a parametric BRDF. Later, Yu et al. [23] propose a method where they measure the incoming illumination from the sun, sky and environment from images and solve for the spatially varying BRDF of an architectural setting. They assume a dichromatic model, with a diffuse and specular component, where the diffuse component varies over the surface and the specular is con-

stant for every surface patch. More recently, Tchou et. al [18] extended the method of Yu et. al by introducing a novel lighting measurement apparatus that can record the high dynamic range of both, sunlit and cloudy environments, using a set of specular and diffuse calibrated spheres. They calculate spatially varying diffuse surface reflectance using an iterative inverse global illumination technique.

3.2. Illumination consistent textures

For applications that do not require an illumination invariant representation of appearance attributes, an illumination consistent texture usually suffices. Some work has been done addressing this problem. Agathos and Fisher [1] have presented a method to create textures that look seamless. Their method works by computing a global 3x3 color correction matrix from a pair of images. They compute this matrix from the overlap region of the two images, and discard pixels in shadow or highlight areas by thresholding. They recolor one of the images applying the color correction matrix, and then they blend the overlapping areas using a weighted average. In computing the color correction matrix they employ an image-based approach that does not make explicit use of the underlying geometry. An extension of this work to multiple overlapping images is presented by Bannai et al. in [4], where they compute multiple color space transformations, first between patches of the same image and then between multiple overlapping regions. Beauchesne and Roy [5] compute illumination consistent textures from the ratio image. They compute the irradiance ratio (which they call “ratio lighting”) for each surface normal in the area of overlap of two images and use it for relighting. To extrapolate for unseen normals they use a filtering mechanism based on a Gaussian kernel. The work on relighting that we present in this paper is similar to their work. We also make use of the irradiance ratio to relight the images. But we extend this technique to outdoor environments, where shadows can be present and must be accounted for.

4. Relighting using irradiance ratio maps

The problem we are solving can be stated as follows. Given,

1. \mathcal{G} the scene geometry, geo-referenced.
2. $\mathcal{I} = \{I_1, I_2, \dots, I_n\}$ a set of photographs of the scene captured under illumination conditions $\mathcal{L} = \{L_1, L_2, \dots, L_n\}$. Some of the images overlap.
3. $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ the set of camera projection matrices that relates \mathcal{G} with \mathcal{I} .

we want to create a textured model of \mathcal{G} as illuminated by one of $L_r \in \mathcal{L}$, i.e. we want to relight all images in \mathcal{I} to

illumination L_r and map them to the model.

In addition, we know at which time each photograph was taken; hence, we also know the sun position at the instant. And we do not know anything about the illumination, except that we are dealing with a naturally lighted scene, where the light comes from the sky and the environment. Shadows will be present and dealt with.

The reflection equation describes how light in the environment is reflected by a surface point \mathbf{x} . For a viewer whose position relative to \mathbf{x} is described by a direction (θ_o, ϕ_o) , the reflected radiance is:

$$B(\mathbf{x}, \theta_o, \phi_o) = \int_{\Omega_i} L(\mathbf{x}, \theta_i, \phi_i) f_r(\mathbf{x}; \theta_i, \phi_i; \theta_o, \phi_o) \cos \theta_i d\omega_i \quad (2)$$

where B is the reflected radiance, L is the incident radiance, i.e. illumination, and f_r is the BRDF. Although we have not explicitly stated so, this equation has an additional dependence on wavelength.

For diffuse surfaces, (2) becomes

$$B(\mathbf{x}) = \rho(\mathbf{x}) \int_{\Omega_i} L(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega_i \quad (3)$$

where ρ is the surface albedo. For simplicity, define the irradiance at \mathbf{x} as

$$E(\mathbf{x}) = \int_{\Omega_i} L(\mathbf{x}, \theta_i, \phi_i) \cos \theta_i d\omega_i \quad (4)$$

Then the reflected radiance equation (3) becomes

$$B(\mathbf{x}) = \rho(\mathbf{x})E(\mathbf{x}) \quad (5)$$

Note that there is still a dependence on the surface position, on both the albedo ρ and the irradiance E . In outdoor environments, the illumination from the sun and sky is far from the scene. If we momentarily ignore cast shadows, and make the assumption of distant illumination, we can reparametrize equation (5) by the surface normal at \mathbf{x} .

$$B(\mathbf{x}) = \rho(\mathbf{x})E(\mathbf{n}(\mathbf{x})) \quad (6)$$

To perform the relighting operation, we will define an irradiance ratio map (IRM). Given both, the target and source illumination environments L_r and L_s , we define the irradiance ratio map parametrized by the surface normal \mathbf{n} as:

$$R_{r,s}(\mathbf{n}) = \frac{E_{L_r}(\mathbf{n})}{E_{L_s}(\mathbf{n})} \quad (7)$$

where $E_L(\mathbf{n})$ is the irradiance at a surface point with normal \mathbf{n} and incident radiance L . Given the IRM R for all surface orientations, it is possible to compute the reflected radiance under the target illumination L_r .

$$B_{L_r}(\mathbf{x}) = \rho(\mathbf{x})E_{L_s}(\mathbf{n}(\mathbf{x}))R(\mathbf{n}(\mathbf{x})) \quad (8)$$

Hence, if we can compute R from the image set \mathcal{I} we can solve the relighting problem.

4.1. Computing the IRM from a set of images

An image is an array of brightness measurements. A single measurement $I(x, y)$ is related to the reflected scene radiance B by the camera response function s . For a point \mathbf{x} in the 3D world, its measured intensity is:

$$I(P(\mathbf{x})) = s(B(\mathbf{x})) \quad (9)$$

where P is the 3D to 2D projection function. By dividing the light spectrum into distinct channels (e.g. RGB) and replication the above equation for each of these, the dependence on wavelength is accounted for. We can now write the IRMs in term of image measurements:

$$R_{r,s}(\mathbf{n}(\mathbf{x})) = \frac{s^{-1}(I_r(P_r(\mathbf{x})))}{s^{-1}(I_s(P_s(\mathbf{x})))} \quad (10)$$

4.2. Extending the IRMs to handle shadows

In an outdoor environment, there are two main sources of light: the sun and the sky, giving rise to two components of incident light, a direct and a diffuse component. Some surfaces will receive light from both, some others will be shadowed and only receive light from the sky. We need to distinguish between two different class of shadows. A surface might be shadowed because its surface normal points away from the sun. We will refer to these shadows as attached shadows. In addition, a surface might also be shadowed because there is an object occluding the sun, giving rise to a cast shadow.

A single IRM can be used to relight in the presence of attached shadows. Cast shadows, however, present a different problem. Irradiance at a surface point that is occluded from the sun will be different from the irradiance at another surface point with the same orientation that is receiving light from the sun. A single IRM will not be enough to handle this situation. Under the assumption that irradiance at all shadowed surface points with the same orientation is similar, it is possible to extend the IRM concept to handle cast shadows. It will be necessary to split the set of surface normals that can receive light from the sun into two sets: a shadowed set and a non shadowed set. Instead of a single IRM, four will be necessary: one for relighting those surface normals that are not in shadow in both images, one for those surface orientations that are in shadow in both images, and two more for the shadowed / non-shadowed cases.

The IRMs are computed from the overlap region between the two images involved in the relighting operation. To relight surfaces in the non-overlap region, the irradiance

ratio for the corresponding orientation is looked up in the appropriate IRM.

5. The relighting pipeline

Our relighting pipeline has three stages: a pre-processing stage, an IRM computation stage, and a relighting stage, as shown in Figure 1. For accurate results, we work directly with the raw images produced by a Nikon D100 camera. These images have 12 bits per channel per pixel, thus giving a higher dynamic range than regular 8 bit images.

5.1. Preprocessing

The input to our relighting procedure is a set of images, the geometric model of the scene, and the corresponding image registration. To build our geometric models we used a Leica HDS 3000 time-of-flight scanner that can take 360 degrees scans. The resulting point cloud (see Figure 3) is tessellated into a triangular mesh. Then, the images and the geometry are registered in the same coordinate system. This can be done manually by picking matching features in the image and the 3D model and computing a projection matrix, or using automatic methods that we have previously developed [17, 19].

Before we can compute the IRMs we need to collect at each vertex of our model two maps: a visibility and a shadow map, as described next.

Visibility map. The visibility map computed at each vertex of the model is a boolean mapping that defines a valid set of images in which the vertex is visible. Three conditions are tested for each image (see also [15]): first, the projection of the vertex must be within the boundaries of the image; second, the the surface normal at the vertex must be pointing towards the camera (i.e., the angle between the optical axis of of the camera and the surface normal vector must be less than $\pi/2$); and, finally, there must be no other surface point in the line connecting the camera and the vertex (i.e. the vertex must be visible from the camera).

Shadow map. The shadow map at each vertex is a boolean map that defines for each image if the vertex was shadowed at the moment the image was taken. To compute this mapping, we record the date and time each image was taken, which together with the latitude and longitude coordinates of the scene are given to a sun position calculator to obtain the position of the sun in the sky. The sun acts as a directional light source. For every mesh vertex, we trace a ray in the direction to the sun. If the ray intersects with the known scene geometry, then the vertex is shadowed.

Image linearization. All images are linearized to remove the effects of the camera response. The camera radiometric calibration is achieved by taking an image of a Macbeth color calibration chart and computing a piecewise linear gain function and a color transformation matrix. We

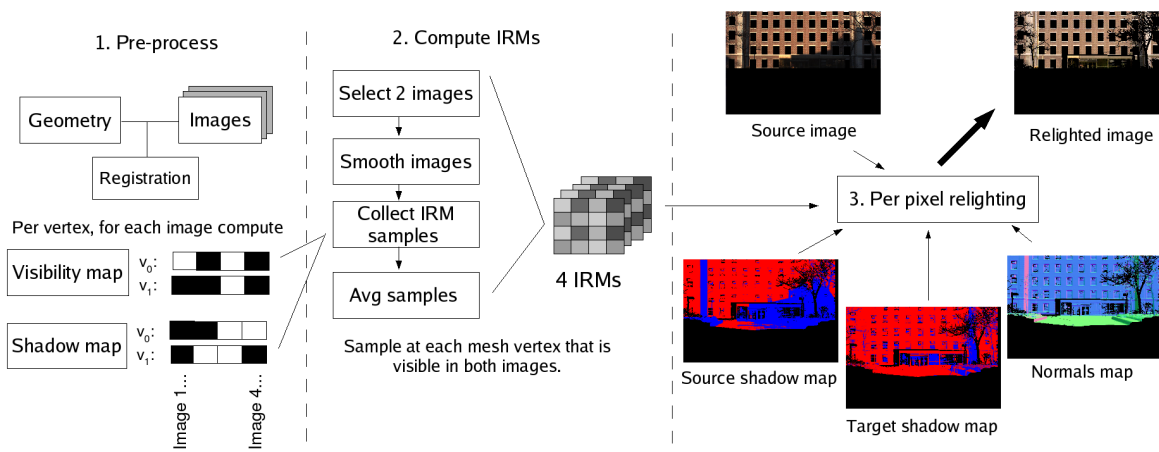


Figure 1. The stages of our relighting pipeline. In the pre-processing stage we compute a shadow and visibility maps for each image at each vertex. We then use this information to collect the IRM samples and compute the 4 IRMs, which are the input to the final stage, where the relighting is performed per pixel. See text for more details.

use the tools provided by the Radiance Lighting Simulation System [20] for this purpose. The resulting linearized images are stored in floating point format and all subsequent operations are performed with floating point precision.

5.2. Computing the IRMs

At this point we select two images, the one to be relighted, or source image, and the image whose illumination we want to match, the target image. To compute the IRMs we iterate over all the vertices of the mesh that are visible in both images (i.e. both visibility bits are set). In addition, we look at the shadow bit for each image to establish which of the four IRMs the irradiance ratio sample will go to. We project the vertices to the images, gather their RGB values and compute the ratio. The IRMs are stored as 2D maps. The surface normal at the vertex is quantized to its corresponding bin in the 2D map and the sample is stored in the IRM. Once all samples are collected, we compute the mean irradiance ratio at each bin in each of the four IRMs.

5.3. Relighting

While we sample a sparse set of points to compute the IRMs, the relighting operation is performed on a per pixel basis, since every pixel in the source image needs to be relighted. Hence, for each pixel in the source image we need to know: its surface normal, its shadow bit for the source image and its shadow bit for the target image. These three values had been previously collected per vertex on the geometric model. To compute these on a per pixel basis, we render the geometry using the computed camera projection

for each image and color encode the required information. First, we create a color encoded rendering of the surface normals. At each vertex, we map its surface normal to an RGB value. The graphics rasterizer does the interpolation between two vertices and fills the intermediate pixel values. In a similar way, we render the shadow maps for the source and target image illumination. A shadowed vertex is given a blue color and a non-shadow vertex is colored red. If a mesh triangle contains blue and red vertices, the colors will be interpolated, allowing also to linearly interpolate between IRMs.

Once these auxiliary images have been generated, we iterate over each pixel in the source image and compute its relighted value in the following manner: first we look at the source and target shadow maps to select the appropriate IRM; then we look up the surface normal, quantize it to its 2D representation and index the selected IRM to obtain the corresponding irradiance ratio; finally, we multiply the source pixel by the irradiance ratio. We ignore those pixels for which we do not have geometry (e.g sky pixels). The resulting image is a floating point high-dynamic range image. In our system, all floating point images are converted to their 8-bit counterpart using the same exposure and color-balance settings.

6. Experimental results

To verify that we can use IRMs to relight an outdoor scene with shadows, we scanned an area of the Columbia Campus (see Figure 3) and acquired several photographs that spanned different illumination conditions. For testing



Figure 2. Two of the images that we used in our experiments. Notice the difference in illumination and shadows (both images are displayed with the same white balance settings). In the relighting test described in the text, the left image is the source and the right image the target.




Figure 3. A point cloud of the model we used in our tests.

purposes (to have ground truth), the photographs overlap almost completely. We masked out windows and trees to exclude them from the relighting operation.

Figure 2 shows two of the images that we used to perform our experimental verification. The left image, which contains shadows cast by a nearby building, was set as the source, and the right image was picked as the target. We ran both images through our relighting pipeline and the results are displayed in Figure 4.

Figure 4 shows the model at different scales both before and after the relighting. In the left column, the left half of the model is textured with the target image and the right half with source image. In the middle column, the left half of the model is textured with the target image and the right half with the relighted version of the source image. With no

relighting, the color transitions between the two textures (marked with an ellipse in Figure 4) are noticeable. After the relighting, the transitions look seamless and the regions that were darkened by shadows (marked with a rectangle in Figure 4) are no longer dark. The rightmost column shows the ground truth; the images in this column are renderings of the model textured with the target image. It can be seen that the relighted image and the ground truth agree well except for a yellow streak in the shadow to non-shadow transition. This yellow streak lies in the penumbra regions of the source image, i.e. the regions where the sun is partially occluded and the shadows are not as dark. Even though we do interpolate between IRMs in these regions, the resulting image still has some artifacts. A potential solution to this problem is discussed below.

7. Discussion and future work

The relighting problem in the context of outdoor environments is a difficult one and little previous work has been done. To the best of our knowledge, our work is the first of its kind that addresses the problem of relighting shadows cast by the sun that does not explicitly measure the incident illumination. Our method is derived from the theory of the image formation process and how light interacts with participating surfaces as opposed to weighted averaging methods. In addition, we have also performed a radiometric calibration of the camera to account for the response function of the sensor, which is ignored in other previous work. Further works need to be done in the following areas:

Penumbra transition. As described earlier, our algorithm does not yet handle penumbras correctly. The size of the penumbra is proportional to the distance to the object that is occluding the sun. Near-by occluding surfaces cast




Figure 4. Results of the relighting operation. The left column shows the model texture-mapped on the left half with the target image and the right half with the source image. The middle column shows the model with the target image on the left half and the relighted source image on the right half. The right column shows the target image texture mapped on the entire model. Row 2 shows two areas of interest. The rectangular area shows how the shadows in the source image have been relighted using the target image. Row 3 is a zoomed image of the elliptical area in row 2 showing the seam from the two images (left) and seamless transition due to the relighting (middle). These images should be viewed in color.

small penumbras, while objects that are far away produce a larger penumbra. Since we know the distance to the occluder, we can use this information to handle the transition smoothly.

Shadow detection. Shadow detection requires knowledge of the complete scene geometry to produce accurate results. In scanning applications one hardly gets the complete geometry and the resulting models usually contain holes. These holes proved to be a problem to us since some shadowed vertices would be marked as non-shadowed (the sun 'leaked' through the hole). We fixed this by manually giving each vertex its correct shadow bit. Further work is required here if full automation is to be achieved.

Non-diffuse surfaces. Our method assumes a diffuse surface reflectance model. There may be imaging conditions which will not meet this assumption. We are currently experimenting using more sophisticated reflectance models

(e.g. Lafortune's model [9]) in the context of IRM relighting.

Extending irradiance ratio maps to unseen surface normals. We compute the IRMs over the overlap region of two images and use them to relight both the overlapping and non-overlapping regions. In the process, we might encounter some surface orientations in the non-overlapping region for which we have not computed an irradiance ratio. Beauchesne and Roy [5] solve this issue by applying a Gaussian filter to the 'ratio lighting' map and computing the missing values. We plan to take a different approach. In outdoor scenes, light comes from the sun and sky. There are empirical and theoretical models available (e.g. Preetham's model [13]) that describe daylight. We would like to use these models to predict from the irradiance ratio samples some parameters for the illumination and hence compute the missing irradiance values for the unseen surface nor-

mals.

Acknowledgements

The authors would like to thank Francesc Moreno-Noguer for his help with the Macbeth charts. This work was funded by NSF ITR grant IIS-0121239 and the Andrew W. Mellon Foundation.

References

- [1] A. Agathos and R. Fisher. Colour texture fusion of multiple range images. In *Proceedings of the 4th International Conference on 3D Digital Imaging and Modeling*, pages 139–146, October 2003.
- [2] P. Allen, S. Feiner, A. Troccoli, H. Benko, E. Ishak, and B. Smith. Seeing into the past: Creating a 3D modeling pipeline for archaeological visualization. In *Proceedings of 2nd International Symposium on 3D Data Processing, Visualization and Transmission*, September 2004.
- [3] P. K. Allen, A. Troccoli, B. Smith, S. Murray, I. Stamos, and M. Leordeanu. New methods for digital modeling of historic sites. *IEEE Comput. Graph. Appl.*, 23(6):32–41, 2003.
- [4] N. Bannai, A. Agathos, and R. B. Fisher. Fusing multiple color images for texturing models. In *Proceedings of the 3D Data Processing, Visualization, and Transmission, 2nd International Symposium on (3DPVT'04)*, pages 558–565. IEEE Computer Society, 2004.
- [5] E. Beauguesne and S. Roy. Automatic relighting of overlapping textures of a 3D model. In *Proceedings of Computer Vision and Pattern Recognition*, 2003.
- [6] F. Bernardini, I. M. Martin, and H. Rushmeier. High-quality texture reconstruction from multiple scans. *IEEE Transactions on Visualization and Computer Graphics*, 7(4):318–332, 2001.
- [7] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432. ACM Press, 2001.
- [8] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: a hybrid geometry-and image-based approach. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 11–20. ACM Press, 1996.
- [9] E. P. F. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126. ACM Press/Addison-Wesley Publishing Co., 1997.
- [10] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.*, 22(2):234–257, 2003.
- [11] R. Love. *Surface Reflection Model Estimation from Naturally Illuminated Image Sequences*. PhD thesis, Leeds, 1997.
- [12] M. Oren and S. K. Nayar. Generalization of Lambert’s reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246. ACM Press, 1994.
- [13] A. J. Preetham, P. Shirley, and B. Smits. A practical analytic model for daylight. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 91–100. ACM Press/Addison-Wesley Publishing Co., 1999.
- [14] K. Pulli, M. Cohen, T. Duchamp, H. Hoppe, L. Shapiro, and W. Stuetzle. View-based rendering: Visualizing real objects from scanned range and color data. In *Rendering Techniques '97 (Proceedings of the Eighth Eurographics Workshop on Rendering)*, pages 23–34, New York, NY, 1997. Springer Wien.
- [15] C. Rocchini, P. Cignoni, C. Montani, and R. Scopigno. Acquiring, stitching and blending diffuse appearance attributes on 3D models. *The Visual Computer*, 18(3):186–204, May 2002.
- [16] I. Stamos and P. K. Allen. Automatic registration of 2-D with 3-D imagery in urban environments. In *Proceedings of the 8th International Conference On Computer Vision (ICCV-01)*, pages 731–737, Los Alamitos, CA, July 9–12 2001. IEEE Computer Society.
- [17] I. Stamos and P. K. Allen. Geometry and texture recovery of scenes of large scale. *Comput. Vis. Image Underst.*, 88(2):94–118, 2002.
- [18] C. Tchou, J. Stumpfel, P. Einarsson, M. Fajardo, and P. Debevec. Unlighting the parthenon. In *Proceedings of the 32nd annual conference on Computer Graphics and Interactive techniques*. ACM Press, August 2004.
- [19] A. Troccoli and P. K. Allen. A shadow based method for image to model registration. In *2nd IEEE Workshop on Image and Video Registration, CVPR'04*, 2004.
- [20] G. Ward and R. Shakespeare. *Rendering with Radiance. Space & Light*, revised edition, 2003.
- [21] G. J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272. ACM Press, 1992.
- [22] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 215–224. ACM Press/Addison-Wesley Publishing Co., 1999.
- [23] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 207–217. ACM Press, 1998.