Towards Automatic Photometric Correction of Casually Illuminated Documents

George V. Landon, Yun Lin, and W. Brent Seales
University of Kentucky
Lexington, KY 40506-0495
{landon,ylin,seales}@netlab.uky.edu

Abstract

Creating uniform lighting for archival-quality document acquisition remains a non-trivial problem. We propose a novel method for automatic photometric correction of non-planar documents by estimating a single, point light-source using a simple light probe. By adding a simple piece of folded white paper with a known 3D surface to a scene, we are able to extract the 3D position of a light source, automatically perform white balance correction, and determine areas of poor illumination. Furthermore, this method is designed with the purpose of adding it to an already implemented document digitization pipeline. To justify our claims, we provide an accuracy analysis of our correction technique using simulated ground-truth data which allows individual sources of error to be determined and compared. These techniques are then applied on real documents that have been acquired using a 3D scanner.

1. Introduction

The need for high-quality acquisition techniques of non-standard documents is becoming a requirement at the forefront of many digital library iniatives. While flatbed scanners have matured to a place where they are the accepted standard for digitization of *standard* documents, documents such as ancient bound manuscripts and scrolls are left without a suitable *mainstream* technology.

While improvements of 3D surface acquisition technology are, allowing the shape of non-planar texts to be captured, the problem of archival-quality texture acquisition still exists. It remains a hard problem for a well-trained photographer to produce near-constant illumination over an aged, and severely wrinkled, manuscript. Moreover, even if suitable lighting is obtained for one portion of a document there is no guarantee that the same lighting configuration will be correct for the next area/page to be acquired.

This work relies on the Shape-from-Shading (SfS) theoretical foundation to estimate light source parameters with the final goal of document correction. The methods presented here combine three fields to produce a final system for document restoration.

ity Office Carpool/Vanpool Program, the use of 27 hybrid electric vehicles by the LFUCG and the use LED lights in Fayette county's 10,088 traffic signals, which use less energy and emit less greenhouse gas.

GreenLex hopes to expand the current bike-lane system before the World Equestrian Games come to Lexington in 2010 to provide environmentally friendly transportation to the event.

UK was not represented on the panel of speakers.

"I think President Todd should have come tonight because he is a strong voice in this community," said forestry junior Drew Combs. "UK is an essential part of the Lexington community. That it was not represented in something as important as this is unfair to the student body and the community."

Two representatives from the Interfaith Alliance of the Bluegrass discussed the link between faith and the environment.

Mark Kline, vice president of the Alliance and panel member, said citizens must speak on behalf of nature because it cannot speak for itself.

"We can't pretend to believe the world will take care of us

Sophomore backup quarterb

DELOTE

Continued from p

Figure 1. (a) Portion of a deteriorated newspaper clipping. (b) Our corrected result.

The first area, Shape-from-Shading, has been one of the major focus area in Computer Vision[10, 22] for many years. There are 4 central parameters involved in this problem: Lighting, surface shape, surface reflectance(albedo), and shading. Normally a set of assumptions are made including: Lambertian surface[13] with constant albedo, directional light source or source at the camera optical center, and an orthogonal projection model [6]. Hertzmann and Seitz [9] proposed an approach that recovers shape and material by looking up a shading-normal table built using a spheric light probe, where distant light and orthographic camera is assumed but the surface is not limited to lambertian. In [2], it has been proved that when the light direction and the Lambertian surface reflectance are unknown, the problem is ill-posed and there is no unique solution. Prados and Faugeras model SfS as a well-posed problem by taking into account the illumination attenuation term [17]: Assuming known light source (at camera optical center), lambertian surface reflectance(constant) and the camera parameters, the shading allows recovery of surface shape from only one image without additional information. They apply

1-4244-1180-7/07/$25.00 ©2007 IEEE

this SfS model to recover cylindric document shape from scanner images and use it for document restoration [16]. Work has also been performed under a point light source assumption. Iwahori et al. [11] performed shape from shading assuming a point light source. Kim et al. [12] also reconstructed shape using photometric stereo and multiple point light source positions. Finally, Frolova et al. [7] demonstrated the use of spherical harmonics to reconstruct depth using point light sources.

However, for standard SfS assumptions the problem remains ill-posed without prior information. Therefore, many groups have added *light probes* that return the problem to a well-defined state. Hara et al.[8] provide a method to estimate a single point light source position and surface reflectance properties from a single image with know scene geometry. This work assumes detectable surface specularity that can be separated from the diffuse surface reflectance.

Using spherical light probes has also provided solutions to light-source calibration. The use of matte spheres that follow a lambertian reflectance properties [25] [23] have also been widely used in light source estimation methods. Powell et al.[15] have introduced a novel calibration object, using 2-3 spheres with known position, to calculate the 3D position of multiple light sources. Also, Takai et al.[20] use the notion of a difference sphere to calculate the position and radiance of multiple light sources. Alldrin et al.[1] use a planar light probe with a multi-layered, transparent medium which differentially absorbs and reflects light. Then Fourier series coefficients are recovered of the incident lighting parameterized over the plane.

Sata et al.[18] have developed robust methods for light estimation using cast shadows of known geometry. While the usefulness of these methods cannot be understated, work in document imaging can often be limited to a library archive environment. These limitations led our work to focus on a light probe that could be manufactured in-field and provide minimal disruption to a digitization pipeline.

Light source estimation is only part of the problem for document correction. The estimated parameters must be applied in a way to provide the best final result.

Document distortion correction techniques have developed from 2D image processing [5] to 3D model manipulation [21, 24, 3, 4, 19]. Due to the fact that the surface shape distortion induces content distortion and shading irregularities in the captured image, techniques have been developed to tackle geometric rectification. A category of the work assumes regular models for the surface shape [21, 24, 4] while more recently a technique based on physically-based simulation [3] has emerged to flatten arbitrarily warped and crinkled documents. Beyond geometric restoration, shading correction has also attracted attention. A few examples of work propose methods to address both geometric and photometric correction of documents of some limited geometric types [24, 4]. Consequently, an algorithm is presented in [19] handles more extensive document types by reducing the shading correction problem to solving a Poisson equation with Dirichlet boundary conditions, which is formulated by relaxing the gradients of the pixels on the illumination edges that correspond to edges in the surface depth map. It is developed to handle documents folded with "hard" creases instead of warped with "soft" distortion, like the document used in our experiment shown in section 4.1. More related to our work, Prados and Faugeras use a Shape-from-Shading algorithm for document restoration [16] as we discussed earlier.

As we have observed in our research, the 3D document acquisition pipeline typically leaves flexibility for improvements to be added between most steps of the digitization process. In this work, we introduce a method for correcting poor lighting conditions during the scanning process. By using a simple folded piece of paper in the acquired scene, we can correct any shading variation caused by a single point light source. It is often a very difficult problem when trying to illuminate a surface uniformly and almost impossible with a single light. This is even more the case when the surface has become distorted from damage, deterioration, and/or age. Furthermore, we can correct white-balance in the final image when a non-white light is used. This entire process adds little overhead and work to the digitization process.

Our method uses the same mathematical foundations of Shape-from-Shading, but in an inverse manner. We know that for a *small* region on a surface, with known albedo, intensity, and normal, the light source direction can be calculated. We take this basic relationship one step further. If we perform this directional light estimation for different patches on the same surface, illuminated by a point light source, can we use the direction for each of these patches to find the 3D location of the light source? We have found that the answer is *yes*. We are able to demonstrate our equations through synthetic data and analyze different sources of error that may affect the accuracy of the calculation. Then we demonstrate our technique on real data that has been acquired with 3D scanners.

2. Method

Derived from Shape-from-Shading theory, we developed an approach to estimate lighting condition, including source position and color/intensity, with a folded white-paper model. Our approach is proposed to deal with both directional light and point light source which is excluded in existing SfS/photometric stereo algorithms. The method is designed to fit into the common scenario of document acquisition and provide automatic digital restoration minimizing extra efforts. Only a simple piece of white folded paper is required and it is added to the acquisition scene serving as a *light probe*, which is scanned to build a model taking advantage of the scanning facilities for the document acquisition.

We assume most documents are originally *flat*. This assumption underlies our restoration framework. Digital images/models of documents are usually distorted because of its non-flat(wrinkled) surface. The distortions can be characterized as photometric distortion and geometric distortion. Using estimated light parameters, we can rectify the photometric distortion, including unwanted tint from the lighting and the non-uniform shading over the surface due to distorted surface geometry.

In parallel with the photometric restoration pipeline, a geometric correction procedure, similar to [3], is applied to the document model to restore the damaged irregular document surface to its original flat status.

Both the light estimation and the restoration procedures are 3D model driven, starting from a 3D model of a document that consists of a geometric mesh and a texture map. Our method integrates with any scanning system that produces a 3D point cloud or surface with a registered texture image.

2.1. Assumptions

A number of assumptions have been applied to Shape-from-Shading and light-source estimation in the literature. However, many of these assumptions have limited real-world applications of the presented methods.

In particular, a common assumption in SfS is the "directional lighting" condition. This assumption easily holds for synthetic experiments, but can be hard to produce with real-world applications. Our method relaxes this strong constraint and provides facilities to estimate point source lighting by making a reasonable assumption:

For a very small surface region, the illumination from a point light source approximates to directional lighting.

Secondly, we assume Lambertian surfaces for both the light calibration target and the document itself. This assumption is easy to approximately satisfy for many commonly used paper types especially historical documents.

In addition, we assume direct illumination i.e. inter-reflections and shadows are neglected. Contingent on this assumption, ambient light must be minimal in the scene; a practical requirement for most document archives. Finally, we assume attenuation of light during transmission in medium (air) can be ignored.

2.2. Light Source Estimation

A piece of white paper is used as a *light probe* to estimate the light source including the position for point sources and direction for directional sources, each in the scanner coordinate frame, and its acquired RGB color. The paper is folded

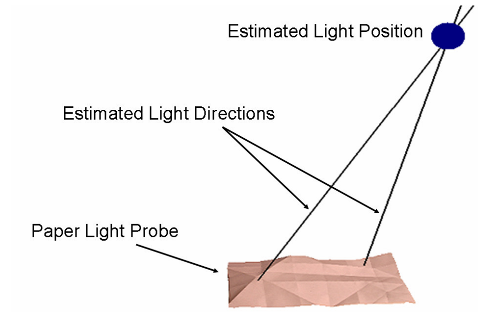

Figure 2. A screenshot of our system with estimated parameters labeled.

such that its surface contains *facets* with non-uniform normals and shading.

2.2.1 Position Estimation

The foundation for our method is the general Irradiance Equation:

$$I(u,v) = I_L \rho(x,y,z) \frac{\vec{L}(x,y,z) \cdot \vec{N}(x,y,z)}{|\vec{L}(x,y,z)||\vec{N}(x,y,z)|} \quad (1)$$

Where I_L is the light intensity/color. ρ denotes the albedo at the shaded point, L is the lighting direction vector and N is the surface normal vector at the point.

Under the perspective projection model, let $X = (x,y,z)$ denote a surface point identified by real-world coordinates. It is projected onto image point $i = (u,v)$:

$$i = \mathbb{P}X \quad (2)$$

where $\mathbb{P}$ is the camera's projection matrix.

Under the point light source condition, we assume that within a "small" region the lighting approximates to directional lighting. We name this kind of region on the lit surface, that is illuminated by approximated directional lighting, as a *patch*. How small a region should be to satisfy the assumption depends on how far the light source is from the lit object. In our experiments, light sources were placed around 1.5 meters and we choose patches with radiuses about 2.5 centimeters. We discuss how the patch size influences the estimation result in section 3.

Suppose N patches are chosen on the calibration target and each patch is illuminated from direction $\vec{L_n} = (lx_n, ly_n, lz_n)$ which is normalized, where $n = 1, 2, ..., N$.

Within patch n, suppose there are M *illumination facets*. An illumination facet is defined as the union of connected polygon faces that share similar normals. Each facet has an intensity I_m and unit normal $\vec{N_m} = (nx_m, ny_m, nz_m)$ where $m = 1, 2, ..., M$. To handle sources of error from real scans, we have added statistical techniques to compute the normal and intensity for each facet by optimally fitting point samples. A Principle Components Analysis (PCA) method is used to compute the normal of the total least-squares fitting plane for a set of vertices within the illumination facet. Also, to compensate for sensor noise in the acquired texture image, we average the intensity values of every texel within an illumination facet.

According to equations 1 and 2, we can construct an irradiance equation for each patch by assuming the calibration target albedo is white (i.e. $\rho = I_{white}$):

$$I_L I_{white}(lx_n nx_m + ly_n ny_m + lz_n nz_m) = I_m \quad (3)$$

Dividing the m^{th} facet by the k^{th} facet yields:

$$\frac{I_m}{I_k} = \frac{lx_n nx_m + ly_n ny_m + lz_n nz_m}{lx_n nx_k + ly_n ny_k + lz_n nz_k} \quad (4)$$

which can be written as:

$$\begin{aligned} & A_{mk} lx_n + B_{mk} ly_n + C_{mk} lz_n = 0 \\ & \qquad A_{mk} = I_m nx_k - I_k nx_m \\ where: & \qquad B_{mk} = I_m ny_k - I_k ny_m \\ & \qquad C_{mk} = I_m nz_k - I_k nz_m \end{aligned} \quad (5)$$

A linear system can be constructed with 3 unknowns: lx_n,ly_n, lz_n for each patch n:

$$\begin{bmatrix} A_{12} & B_{12} & C_{12} \\ A_{23} & B_{23} & C_{23} \\ ... & ... & ... \\ A_{(M-1)M} & B_{(M-1)M} & C_{(M-1)M} \end{bmatrix} \begin{bmatrix} lx_n \\ ly_n \\ lz_n \end{bmatrix} = 0 \quad (6)$$

At minimum, 4 facets are needed to solve the lighting direction for each patch. Usually more facets are used to overcome image noise and other factors. Singular Value Decomposition (SVD) is applied on the coefficient matrix to find the null-space of the system and obtain the optimal least-squares solution.

For directional light, 1 patch is enough to obtain the light direction. For a point light source, after having obtained the lighting directions for at least 2 patches, we calculate the intersection of the light rays associated with the patches to obtain the point light source position S, which can be defined through the parametric representation of each light ray:

$$S = C_n + \vec{L_n} t_n \;\; n = 1, 2, ..., N \quad (7)$$

which leads to the linear system:

$$\begin{bmatrix} 1 & -\vec{L_1} & 0 & ... & 0 \\ 1 & 0 & -\vec{L_2} & ... & 0 \\ ... & ... & ... & ... & ... \\ 1 & 0 & 0 & ... & -\vec{L_N} \end{bmatrix} \begin{bmatrix} S \\ t_1 \\ t_2 \\ ... \\ t_N \end{bmatrix} = \begin{bmatrix} C_1 \\ C_2 \\ ... \\ C_N \end{bmatrix} \quad (8)$$

where C_n is the center of patch n. The solution can be found by multiplying the right side matrix of the equation by the pseudo-inverse of the coefficient matrix.

In practice, we apply a RANSAC algorithm to overcome bad patch estimations due to un-modeled reflectance properties and a lack of surface normal variations.

2.2.2 Color Estimation

With estimated light source position $S = (sx, sy, sz)$, each illumination facet m can vote for a light color estimation by solving I_L from equation 3:

$$I_{Lm} = \frac{I_m}{I_{white}(\frac{\overrightarrow{S - C_m}}{|S - C_m|} \cdot \overrightarrow{N_m})} \quad (9)$$

where C_m is the center of the facet. We average the voted colors as the final estimation. The operations are applied independently on the R, G, and B color channels. This allows us to solve for the captured light source color.

2.3. Document Restoration

A complete document restoration is achieved by simultaneously correcting the distortion from the two components of a document model: the geometric mesh and the texture map. We implement geometric correction using the techniques presented in [3, 19]. A surface mesh acquired from a high-resolution 3D scanner, modeled as a mass-spring system deforms realistically when simulation is governed by physical laws. A ”gravity” force is simulated to drive the surface to collide with and rest on the flat ground. The whole system strives to minimize its accumulated potential energy of the springs. The technique is able to restore an arbitrary crumpled document surface and produce an approximate conformal mapping that parameterizes the three dimensional surface on a planar domain [14].

Many documents, especially aged documents, are close to Lambertian surfaces, where the reflectance varies directly proportional to a scale factor f that equals to the cosine of the angle θ between the lighting direction and the normal of the surface at that point given by $f = cos\theta$. With an estimated point light source position, using our approach described in section 2.2.1, it is straightforward to compensate for the non-uniform illumination over the surface: the intensity value of each image pixel is divided by $cos\theta$. The resultant image is equivalent to one captured for a document with flat surface. Moreover, the scene lighting is not perfectly white in most cases and imaging devices often add skew to the color of a scene. With the captured light color estimated in section 2.2.2, we are able to remove the tinting effect caused by the illumination and restore the document image to its original albedo by updating the scale factor f as $f = I_L cos\theta$. When completed, the final output is a document image without photometric and geometric distortion.

3. Empirical Analysis

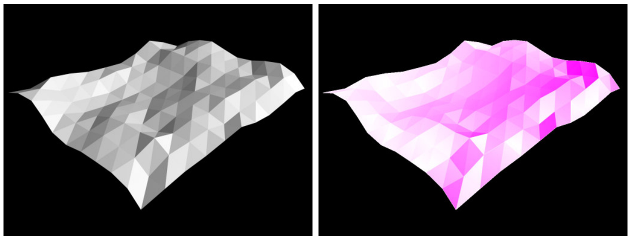

Figure 3. (a) The ground-truth shaded surface. (b) The difference image of the ground-truth shaded surface and shaded surface with vertex error added ($\sigma = 1.2$). Note: White (diff$=0$) and Red (diff$=255$).

We present analysis of various aspects of our method with ground-truth data. Using synthetic data, we are able to control most sources of error. Therefore, errors can be analyzed independently by simulation in a controlled environment. After quantifying the error, we show results from an actual scan performed with a 3D scanner.

3.1. Surface Perturbations

All readily available 3D scanning technologies introduce some form noise in the acquired surface. High-frequency noise, in particular, is very common in most scanning technologies. These small variations in the surface add great amounts of error when triangulating a 3D point cloud. This error propagates directly to the surface normal for each polygon which is used directly when estimating the light-source position. To determine how great the effect of high-

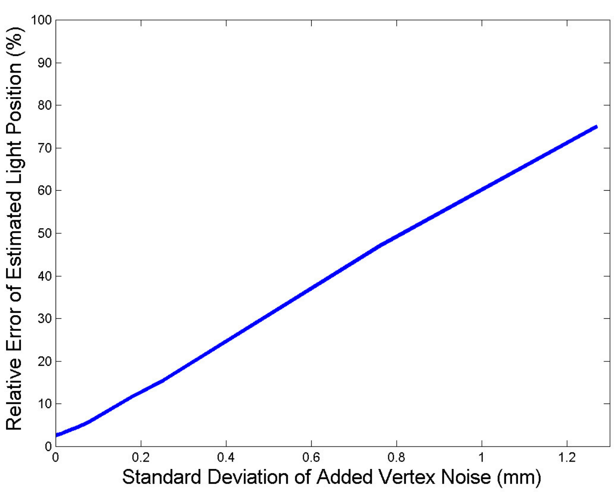

Figure 4. Demonstration of the error introduced by random noise in the 3D points of a surface.

frequency surface noise is on our final estimation, we performed experimental calculations directly on ground-truth data. Using a random number generator following a uniform distribution, we added varying levels of random error to each dimension of the 3D vertices.

Figure 4 shows the result of applying the surface perturbations. The plot shows that error follows a linear pattern.

Figure 3 shows the difference of the ground-truth shaded image with the shaded image of an estimated light position with added error.

3.2. Number of Patches Selected

During our experiments, we have determined that the number of patches selected also plays a role in the accuracy of the final light position result. To test this result, we

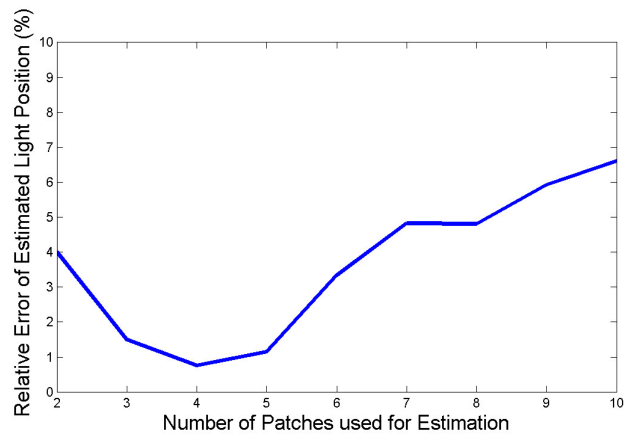

Figure 5. Demonstration of the accuracy produced with varying numbers of patches selected.

selected an increasing number of patches with 5 facets each and calculated the light position with each increase. We started at the minimum of 2 patches and added up to 8 more patches. The two most notable features of this plot are the global minimum when there are 3, 4, or 5 patches and, once more than 5 patches are used, error begins to greatly reduce the overall accuracy of the light estimation.

3.3. Number of Facets per Patch

Another parameter in our algorithm is the number of facets selected per patch. The underlying equations require atleast 4 facets to solve the set of equations. However, as can be seen in Figure 6, once more than 8 facets are chosen per patch, error greatly reduces the estimation. This validates our assumption that a small area on the surface can be estimated with directional light, but when this area becomes to large this assumption fails. Furthermore, depending on the selection of the facets, when only using a small number, their variation in intensity and normals plays a very important role. For example, it should be possible to select 4 facets with similar intensities and normals. This case is under-constrained and will lead to an invalid solution.

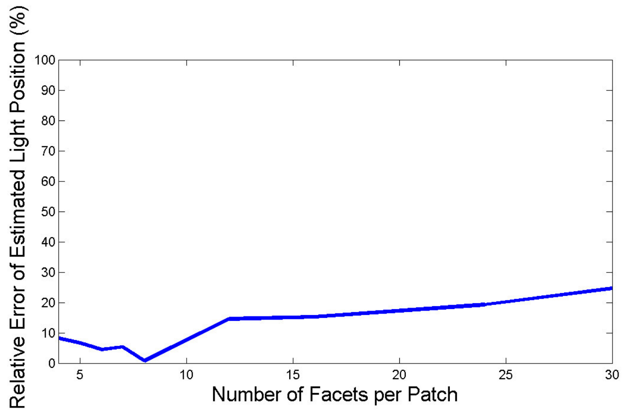

Figure 6. Demonstration of the accuracy produced by increasing the total number of facets for every patch.

3.4. Other Sources of Error

Another source of error can be the placement of the patches themselves. Since the algorithm relies on triangulation of directional light rays, a user must keep in mind that patches to close together might produce parallel or near-parallel rays. This issue requires the simple solution of selecting patches with enough relative distance between them to produce nicely-behaving intersecting rays.

On a different note, there is a different source of error introduced by the texture acquisition itself. Digital cameras introduce another source of error when imaging a surface. Gaussian noise is an intrinsic property of current image sensor technologies which becomes even more apparent when dealing with areas of low illumination in a scene.

3.5. Real-world Analysis

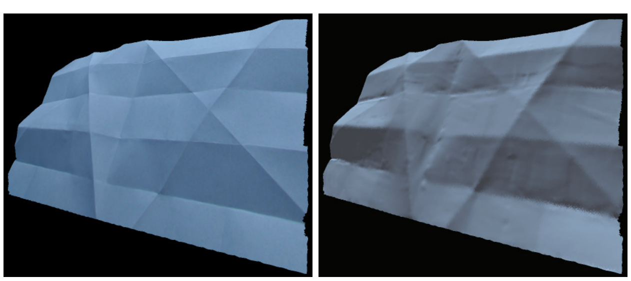

Figure 7. A comparison between acquired texture image and shaded rendered image with estimated light position. (a) shows texture image. (b) shows shaded image with light estimated.

Since performing exact measurements of actual light sources is very difficult, providing numerical analysis of real-world results becomes much harder. However, demonstrating the performance of our algorithm by visual observation serves as a sufficient example that our results are as expected. In Figure 7, the 3D surface of an acquired calibration target is shown. For reference, the acquired texture image, used for calibration, is shown in Figure 7(a). The estimated light source position and color is then calculated for this data set and the rendered estimation is shown in Figure 7(b). As can be seen in Figure 7, the shaded surface looks consistent with the acquired texture image.

3.6. Effect of Error on Photometric Correction

The accuracy of the photometric correction with respect to the original document is directly related to the accuracy of the light source position and color estimation. In Figure 8, a direct comparison is made between a ground-truth document image and the photometrically corrected image for varying degrees of light source position accuracy. For reference to the previous plots, we move the Y-axis, error of light position, to the X-axis and show the error between the ground truth and the photometrically corrected image for each light source position on the Y-axis.

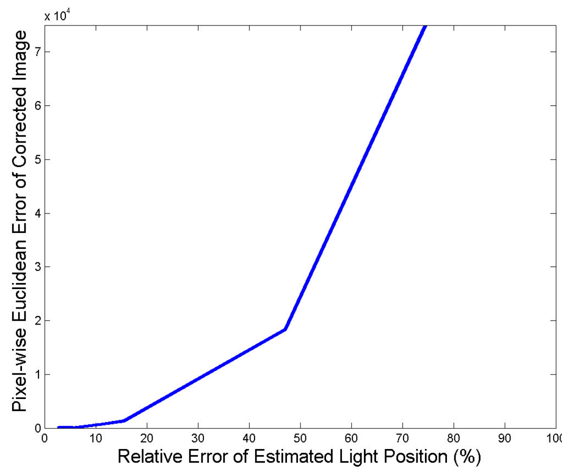

Figure 8. Overall photometric correction accuracy dependent on light source estimation.

4. Results

We have implemented the algorithms presented in section 2 and applied them on synthetic and real-world document restoration.

4.1. Synthetic Results

Using Autodesk 3D Studio Max, we generated a virtual scene that included a distorted surface with a document texture map, a distorted white surface to serve as a light probe, a colored point-light source that illuminates the scene, and a camera. We created the surface distortions using the built-in noise function and adding a preservation constraint that kept the polygon edge lengths close to the original edges. Figure 9(a) shows the image captured by the camera. The geometric distortions cause the non-uniform illumination over the

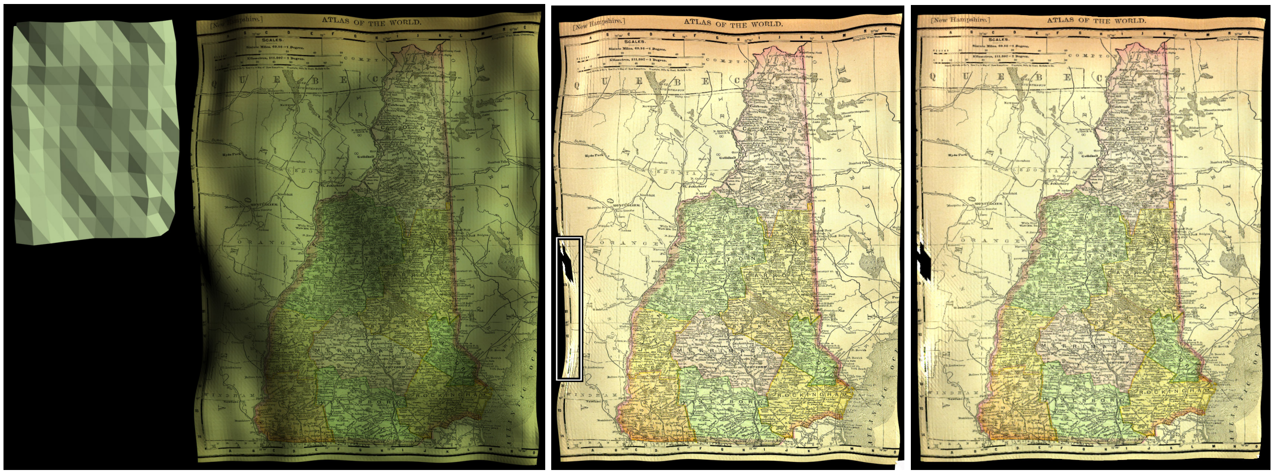

Figure 9. (a): The synthetic input image of a distorted document and paper light probe. (b): The photometrically corrected document image using the proposed method. (c): The final output showing geometric and photometric distortion removal.

surface and the distorted lines. Also, due to the influence of the light color, the captured document image deviates from its original albedo.

We applied the method described in section 2.2 to obtain an estimation for the light position and its color. Using the estimated values, photometric correction was applied and produced a document image free from distortion when combined with the geometric correction. The estimated light position is (14.1074, -35.5189, 40.0826) versus ground truth (14.243; -36.644; 40.965) and estimated light color (RGB) is (188.7, 213.69, 155.55) versus ground truth (189,213,155). Figure 9 shows the entire restoration process. Figure 9(a) shows the input image, (b) shows photometric correction, and (c) shows the final output after geometric correction. The document content looks significantly better than the input image: lines are straightened, illumination is uniform and albedos are undistorted. In Figure 9(b), the rectangle marks the region that failed in restoration due to shadows in the input image. However, our system can notify the user of these failures to facilitate proper action during digitization.

4.2. Real-world Results

In Figure 10(a), we show an image of a wrinkled newspaper. These non-uniform variations in the surface create a very difficult task when trying to generate constant illumination across the surface. The result of our photometric correction and geometric un-warp is shown in Figure 10(c). In Figure 11(a), we show the image of a heavily worn book cover. This surface is particularly interesting because of water-damage and folding that have occurred to the surface. The result of our photometric correction is shown in Figure 11(b). However, it should be noted that the 3D scanner we used did not have high enough resolution to model all of the surface variations. Therefore, some of the dimples and folds are not completely corrected. However, this is a limitation of our scanning device which did not accurately model the surface and is not a failure of our method.

Figure 10. (a) Example of damaged newspaper fragment. (b) The photometric and geometric correction result. (The highlighted square marks the region shown in Figure 1.)

5. Conclusion

In this work, we have presented a novel technique for the photometric correction of distorted documents that can be included in most current restoration systems, as we have demonstrated. This system robustly estimates the position and color of a point-light source using a new type of light probe that reduces many of the requirements needed for calibration. We have demonstrated particular sources of error that may have adverse affects on the results. Moreover, we have produced promising results for both synthetic and real-world data.

References

[1] N. G. Alldrin and D. J. Kriegman. A planar light probe. In *CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*,

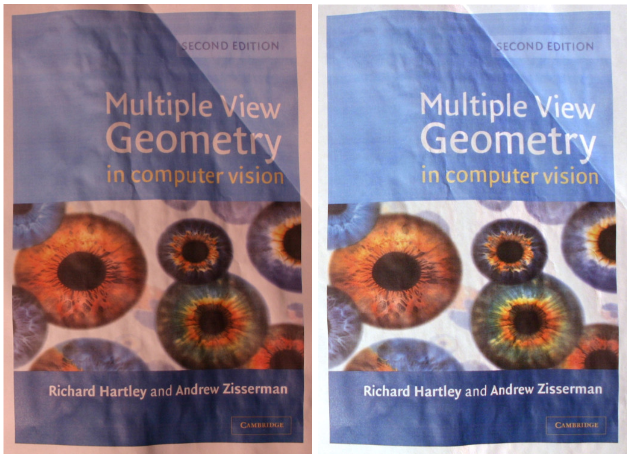

Figure 11. (a) Example of a well-used book cover with water damage and folding. (b) Our photometrically corrected result.

pages 2324–2330, Washington, DC, USA, 2006. IEEE Computer Society.

[2] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *Int. J. Comput. Vision*, 35(1):33–44, 1999.

[3] M. Brown and W. B. Seales. Image restoration of arbitrarily warped documents. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1295–1306, October 2004.

[4] M. Brown and Y. Tsoi. Geometric and shading correction for images of printed materials using boundary. *IEEE Trans. Image Proc.*, 15(6):1544–1554, June 2006.

[5] A. Chaudhuri and S. Chaudhuri. Robust detection of skew in document images. *IEEE Trans. on Image Proc.*, 6(2):344–349, 1997.

[6] M. S. Drew. Photometric stereo without multiple images. In B. E. Rogowitz and T. N. Pappas, editors, *Proc. SPIE Vol. 3016, p. 369-380, Human Vision and Electronic Imaging II*, pages 369–380, June 1997.

[7] D. Frolova, D. Simakov, and R. Basri. Accuracy of spherical harmonics approximation for images of lambertian objects under far and near lighting. In *ECCV 2004*, pages 574–587, 2004.

[8] K. Hara, K. Nishino, and K. Ikeuchi. Light source position and reflectance estimation from a single view without the distant illumination assumption. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):493–505, 2005.

[9] A. Hertzmann and S. M. Seitz. Shape and materials by example: A photometric stereo approach. volume 01, page 533, Los Alamitos, CA, USA, 2003. IEEE Computer Society.

[10] B. Horn. Shape from shading: a method for obtaining the shape of a smooth opaque object from one view. *MIT Project Mac Technical Report TR-79*, 1970.

[11] Y. Iwahori, H. Sugie, and N. Ishii. Reconstructing shape from shading images under point light source illumination. In *IEEE ICPR 90*, volume I, pages 83–87, 1990.

[12] B. Kim and P. Burger. Depth and shape from shading using the photometric stereo method. *CVGIP: Image Underst.*, 54(3):416–427, 1991.

[13] J. Lambert. Photometria sive de mensura et gradibus luminus colorum et umbrae. *Eberhard Klett*, 1760.

[14] Y. Lin and W. B. Seales. Opaque document imaging: Building images of inaccessible texts. In *Proceedings of the 10th IEEE International Conference on Computer Vision*, volume 1, pages 662–669, October 2005.

[15] M. W. Powell, S. Sarkar, and D. Goldgof. A simple strategy for calibrating the geometry of light sources. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(9):1022–1027, 2001.

[16] E. Prados, F. Camilli, and O. Faugeras. A unifying and rigorous shape from shading method adapted to realistic data and applications. *Journal of Mathematical Imaging and Vision*, 25(3):307–328, October 2006.

[17] E. Prados and O. Faugeras. Shape from shading: a well-posed problem? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'05), San Diego, California*, volume II, pages 870–877. IEEE, June 2005.

[18] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(3):290–300, 2003.

[19] M. Sun, R. Yang, Y. Lin, G. V. Landon, W. B. Seales, and M. S. Brown. Geometric and photometric restoration of distorted documents. In *Proceedings of the 10th IEEE International Conference on Computer Vision*, volume 2, pages 1117– 1123, October 2005.

[20] T. Takai, K. Niinuma, A. Maki, and T. Matsuyama. Difference sphere: An approach to near light source estimation. *cvpr*, 01:98–105, 2004.

[21] T. Wada, H. Ukida, and T. Matsuyama. Shape from shading with interreflections under a proximal light source: Distortion-free copying of an unfoldedbook. *Int. J. Comput. Vision*, 24(2):125–135, 1997.

[22] R. J. Woodham. *Photometric method for determining surface orientation from multiple images*. MIT Press, Cambridge, MA, USA, 1989.

[23] Y. Zhang and Y.-H. Yang. Multiple illuminant direction detection with application to image synthesis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(8):915–920, 2001.

[24] Z. Zhang, C. L. Tan, and L. Fan. Estimation of 3d shape of warped document surface for image restoration. In *ICPR '04: Proceedings of the Pattern Recognition, 17th International Conference on (ICPR'04) Volume 1*, pages 486–489, Washington, DC, USA, 2004. IEEE Computer Society.

[25] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(7):680–702, 1991.