

# Restoring 2D Content from Distorted Documents

Michael S. Brown, *Member, IEEE*, Mingxuan Sun, Ruigang Yang, *Member, IEEE*,  
Lin Yun, *Student Member, IEEE*, and W. Brent Seales, *Member, IEEE*

**Abstract**—This paper presents a framework to restore the 2D content printed on documents in the presence of geometric distortion and nonuniform illumination. Compared with text-based document imaging approaches that correct distortion to a level necessary to obtain sufficiently readable text or to facilitate optical character recognition (OCR), our work targets nontextual documents where the original printed content is desired. To achieve this goal, our framework acquires a 3D scan of the document's surface together with a high-resolution image. Conformal mapping is used to rectify geometric distortion by mapping the 3D surface back to a plane while minimizing angular distortion. This conformal “deskewing” assumes no parametric model of the document's surface and is suitable for arbitrary distortions. Illumination correction is performed by using the 3D shape to distinguish content gradient edges from illumination gradient edges in the high-resolution image. Integration is performed using only the content edges to obtain a *reflectance image* with significantly less illumination artifacts. This approach makes no assumptions about light sources and their positions. The results from the geometric and photometric correction are combined to produce the final output.

**Index Terms**—Document restoration, geometric correction, shading correction, photometric correction, conformal mapping, document processing.

## 1 INTRODUCTION

THE digitization of documents is a crucial first step in the process of building an archive to support preservation, access, and continued scholarship. Although printed documents have traditionally been digitized using flatbed scanner technology, an increasing trend is to use cameras [30]. Camera-based imaging allows fast noncontact imaging of printed materials and can accommodate those materials not suitable for flatbed scanners, such as large and bulky items. One consequence, however, is that camera-based imaging does little to guarantee that items are flat when imaged. As a result, unwanted distortion may be present in the final image and will need to be corrected by postprocessing algorithms.

The degree to which distortion correction must be performed depends on the goal of the digitization. The vast majority of imaged items are text documents where importance is placed on extracting the text printed within. In such cases, distortion must be corrected only to the extent of being able to extract text for tasks such as reading or subsequent optical character recognition (OCR). For such purposes, the final output is most often bitonal images.

There are, however, a number of imaged items where the original printed content itself is desired. Examples include items such as artwork, maps, photographs, or primary source materials such as handwritten documents and historical records. For such items, the goal of the digitization is to produce digital facsimiles that represent the original printed content as faithfully as possible. The work presented in this paper targets this latter class of materials from which the original 2D content is desired.

The problem then is straightforward: although printed content is 2D, the material itself after printing does not necessarily remain planar. This gives way to two types of distortion that must be corrected after imaging: 1) geometric distortion and 2) nonuniform illumination distortion. Printed materials may not be flat for several reasons, including folding and bending, damage and decay, or intentionally by design. Fig. 1 shows an example of a document that suffers from arbitrary folding. Image intensities comprising the printed content of the document are distorted by the document's geometry and by shading artifacts from nonuniform illumination (also attributed to the geometry). Fig. 1b shows that correcting geometric distortion alone is not sufficient as the document still appears distorted due to the shading cues. To fully recover the printed 2D content, both geometry and illumination must be addressed, as shown in Fig. 1c.

Geometric distortion correction of documents has received a great deal of attention. Depending on the task, techniques ranging from simple 2D planar deskewing to distortion correction for arbitrary geometries have been developed (for example, [25], [34], [47]). On the other hand, the problem of correcting shading in the context of document restoration, beyond the purpose of simple binarization for text documents, has received significantly less attention. Nevertheless, these are companion problems: shape distortion induces shading irregularities that are difficult to remove at image acquisition. Some of the few

- M.S. Brown is with the School of Computer Engineering, Nanyang Technological University, Blk N4, 2A-32, Nanyang Avenue, Singapore 639798, Republic of Singapore. E-mail: msbrown@ntu.edu.sg.
- M. Sun is with the GVI Center, Georgia Institute of Technology, 85 Fifth St. NW, Atlanta, GA 30332-0760. E-mail: cynthia@cc.gatech.edu.
- R. Yang is with the Computer Science Department, University of Kentucky, 232 Hardymon Building, Lexington, KY 40506. E-mail: ryang@cs.uky.edu.
- Y. Lin and W.B. Seales are with the Laboratory for Advanced Networking, University of Kentucky, Hardymon Building, 2nd Floor, 301 Rose Street, Lexington, KY 40506-0495. E-mail: {ylin, seales}@netlab.uky.edu.

Manuscript received 7 Apr. 2006; revised 9 Oct. 2006; accepted 16 Jan. 2007; published online 22 Feb. 2007.

Recommended for acceptance by J. Hu.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0271-0406. Digital Object Identifier no. 10.1109/TPAMI.2007.1118.



Fig. 1. (a) Original image exhibiting geometric and illumination distortion. (b) Geometry corrected. (c) Geometry and illumination corrected.

examples of work addressing both shape and shading restoration [8], [47] assume restrictive geometric models that limits their utility in general cases.

**Our contribution.** We present a system to rectify both the geometry and illumination of distorted documents. The type of 3D document shape we handle can be *arbitrary*; it is not necessary for the shape to fit any particular parameterized model. Our approach focuses on document imaging tasks that target non-text-based materials. Therefore, assumptions about the document's printed content cannot be exploited. Furthermore, the illumination correction cannot be handled by simple binarization or image filtering.

To achieve our goal, we have designed a system that acquires a 3D reconstruction of the document surface together with a registered high-resolution image. Two-dimensional conformal parameterization of the document's shape is computed in order to obtain a mapping that removes shape-based distortion. The 3D information is also used to address nonuniform illumination by allowing image intensity changes (that is, the image gradient) to be reliably classified into two sets: *illumination changes* and *reflectance changes*. This classification scheme allows the gradient field to be modified by removing all illumination edges caused by the distorted surface normal, which is a direct result of variations in 3D shape. This modified image gradient field allows the computation of a *corrected* image with shading artifacts significantly lessened.

The approaches of our system are cooperative in nature, where the 3D shape and input image are used to correct for both geometry and shading, and the results from these procedures are combined to generate the final restored 2D document. The approaches presented in this paper have several advantages over previous work in document restoration. For geometric correction, our conformal mapping approach is significantly faster than energy minimization approaches and is solved as a linear system. With regards to illumination correction, our approach does not need to explicitly compute the number of light sources in the scene, nor their positional information, which is necessary for related simulation-based techniques that also utilize 3D shape. From a system standpoint, our contribution is a complete pipeline for geometric and photometric restoration of arbitrarily distorted documents.

Shorter versions of this work have appeared in [9], [35]. The work presented in [9] demonstrated how conformal parameterization can be used for geometric correction,

whereas the work presented in [35] introduced the technique to correct shading artifacts. We expand our previous work to provide more clarity to the algorithms and more thorough experimentation, including synthetic examples for quantitative evaluation.

The remainder of this paper is organized as follows: Section 2 presents related work, Section 3 discusses the acquisition system for obtaining the initial 3D data and registered high-resolution image, Section 4 describes the conformal parameterization for geometric correction, and Section 5 describes the shading correction technique. Experimental results are presented in Section 6. Sections 7 and 8 provide a closing discussion and conclusion.

## 2 RELATED WORK

Related work is divided into two categories based on either a geometric or a photometric emphasis. Work involving geometric correction is discussed in the context of document restoration and mesh parameterizations.

### 2.1 Geometric Correction

#### 2.1.1 Document Restoration

Distortion correction algorithms have traditionally focused on planar skew found in flatbed-imaged items where the imaged content is not aligned with the image axis. These approaches use various techniques to compute affine or planar transforms to deskew the image (for example, see [24], [25]). Another type of distortion addressed is the curvature distortion near the spine of a book. Although this distortion is caused by the material's nonplanar shape, the restrictive and predictable nature of this deformation makes it relatively easy to correct using cylindrical or cylindrical-like models [11], [10], [24], [25], [36], [39], [42], [45]. In some cases, even effects from the text degradation due the scanner's sensor are taken into account in the restoration process [24], [25].

Another class of algorithms exploit some a priori knowledge about the imaged materials. Most of these algorithms exploit the assumption that imaged content is of text that should appear as straight lines and, therefore, try to rectify distortion based on how the content deviates from this straight-line assumption [29], [41], [46]. Other algorithms consider the a priori known shape of the document itself, for example, the boundaries of the document should be rectangular. Given the distorted document's image

boundaries, the imaged content can be rectified back to a rectangle using various algorithms such as constraints on applicable surfaces or boundary interpolation [8], [21].

Techniques have been developed to address arbitrary distortion that make no assumptions to the document's printed content or boundary. These approaches begin by [6], [7], [34]. In these approaches, 3D structured-light systems are used with a high-resolution camera to capture an image of the document along with an approximation of the document's surface as a 3D triangulated mesh. For sufficiently textured documents (such as text-based documents), two stereo cameras can be used instead of structured lighting to reconstruct the 3D surface [40].

The idea now is to "flatten" the 3D surface textured with the image of the document to generate a corrected representation. Although the 3D surface of a document should be *developable* to a plane (that is, mapped to a plane without distortion), it cannot be "flattened" by simple 2D retiling of the surface triangles [34]. This is because the 3D mesh represents only an approximation of the true surface, subject to 3D reconstruction errors and undersampling. Thus, the reconstructed surface is not truly isometric with a plane. This results in the need to deform or overlap some of the triangles in the flattening process. Overlapping triangles would result in undesirable artifacts in the restored image and is unacceptable in a restorative context. As a result, *some* notion of deformation (that is, distortion) must be allowed in the 3D to 2D mapping process.

Initial "flattening" techniques used relaxation methods that maintained surface geodesics by minimizing distances between the 3D vertices as they are pushed flat to a plane [6], [7], [34]. This can be performed by associating some energy with changes in adjacent vertices distances, such as spring energy. This transforms the flattening problem into one of energy minimization. Although these approaches are successful, they suffer from being computationally slow. Brown and Seales [6] state that their approach takes roughly one minute to converge on a regularly sampled mesh with  $46 \times 46$  vertices. Pilu states [34] that over 1,000 iterations on a grid of  $20 \times 15$  3D points are needed. Although recent work [12] has improved the computation time by changing integration styles and distance energy used in the minimization process, the algorithms are still iterative in nature.

In this paper, we address this flattening problem using conformal mapping. A conformal map is a 2D parameterization of a 3D surface such that angles are preserved. Developable surfaces are completely conformal, which is intuitive by their definition since no shearing or stretching is necessary to map them to a plane. For general surfaces (including nearly developable surfaces), a conformal map can be computed that minimizes angular distortion. Note that this is different from the distance preserving metrics initially presented in [6] and [34]. Angular preservation can be considered a weaker constraint than distance preserving, because they allow for local changes in scale. However, as we demonstrate in this paper, for the purpose of document restoration, this conformal constraint can restore images of distorted documents to within a single pixel of their "flattened" representation and is on a par with distance preserving techniques. In addition, the conformal parameterization can be computed quickly by solving a linear system.

In the context of document correction, the closest to our work is that proposed in [1]. In this work, triangles making up

the 3D mesh are mapped to a plane one by one, using the results of the previous triangle's mapping to guide the subsequent triangle's mapping. This one-by-one tiling resulted in relatively poor results and provided no mechanism to globally distribute the errors over the whole mesh.

### 2.1.2 Mesh Parameterization

Because our distortion correction algorithm operates on a 3D mesh representing the document's surface, related work can be found in the areas of computer graphics and geometric modeling, where 2D parameterization of 3D meshes is of great interest for remeshing, surface fitting, and computing texture coordinates. Approaches in this area vary by the distortion metric used and the underlying approaches employed. For a good overview of various techniques, we refer the reader to [13].

Most relevant to our approach is that proposed in [18] that parameterizes a 3D point cloud onto a 2D plane. This work is significant because the algorithm could be formulated as a linear system. This approach, however, has one caveat that the initial boundary of the 3D mesh must be specified on the 2D plane. For document correction, this would require that the shape of the boundaries of the document be known in advance. Although this may be reasonable if the document's boundary can always be guaranteed to be rectangular, in practice, however, this is not the case. Often, only a portion of the document surface is extracted for flattening (for example, see Fig. 5). The boundary of this extracted patch is not known before restoration. Furthermore, some printed material's boundaries are not flat by design (for example, see Fig. 12). Thus, a priori knowledge about the resulting print materials boundary is a constraint we would like to avoid.

Recently, Desburn et al. [16] and Lévy et al. [28] presented conformal parameterization techniques that could handle arbitrary boundaries and could be cast as linear systems and solved efficiently. Although these techniques were targeted at a different set of applications such as mesh resampling and parameterization for texture mapping, their basic idea served as the impetus for our work, and we have adopted them for use in document restoration.

## 2.2 Photometric Correction

Correcting shading distortion in documents, outside simple image thresholding used on black-and-white text-based documents, has received significantly less attention. Outside the area of document processing, however, one finds a rich set of approaches to shading correction. The intrinsic image [3], for example, introduces the idea of decomposing an image into two intrinsic images: a reflectance image and an illumination image. With this intrinsic decomposition, the reflectance image is decoupled from the scene's illumination. Recovering these two intrinsic images from a single image, however, is an ill-posed problem as the number of unknowns is twice the number of equations [43]. Nevertheless, progress has been made toward achieving this decomposition under certain conditions. Many techniques (for example, [5], [20], [22], [32]) are based on the *Retinex* theory [27], which assumes that the gradients along reflectance changes have much larger magnitudes than those caused by illumination. These techniques work well for slow-varying illumination changes but cannot handle sharp illumination edges (as seen in Fig. 1). Recent algorithms have used machine learning techniques to separate illumination and reflectance [4], [37].

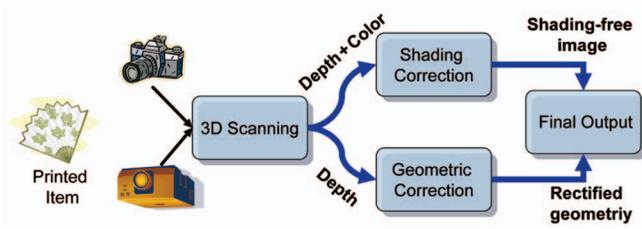


Fig. 2. A system diagram of the overall document restoration process.

Although they can deal with more complex situations, the result depends largely on the training data set.

The above approaches typically do not require 3D information but make some strong assumptions about the scene. In a different direction, research in computer graphics uses 3D information to recover arbitrary textures (reflectance) [15], [19], [31], [44]. These approaches use physically-based light transport simulation to remove the effect of existing lighting conditions. A priori information of the lighting conditions is typically required for good results.

For our approach, our shading correction exploits the presence of the 3D geometry, necessary for geometric correction. The 3D geometry is used to classify edges as illumination edges or reflectance edges. This will be used to remove the illumination edges to help rectify the overall illumination of the 2D content.

### 3 SYSTEM OVERVIEW

Fig. 2 shows a diagram of our overall framework. A 3D reconstruction and high-resolution image of a document are acquired using a structured light system. Geometric and shading corrections are performed separately. The two results are then combined, using the mapping for the geometric correction with the corrected illumination image to produce the final output. Sections 3, 4, and 5 detail our acquisition setup and the algorithms used to correct the geometry and shading distortion.

Our 3D scanning system (shown in Fig. 3a) uses two digital video cameras, a high-resolution still digital camera, and a laser light source. The video cameras and the high-resolution digital camera are calibrated using a physical calibration pattern [38]. Since the laser is used only as a lighting device for creating correspondences, it is not necessary for it to be calibrated. The laser is mounted on a PC-controlled pan/tilt unit and is fitted with a slot lens to form the beam into a thin plane. The pan/tilt control allows the laser plane to be swept across the document. The laser stripe generator, with 5-mW power output at 635-nm wavelength, produces a 1-mm-wide stripe on the document surface. The laser plane is imaged as a stripe by the two video cameras at a resolution of  $640 \times 480$ . The stripe is localized to subpixel accuracy by fitting a Gaussian curve to the intensity profile. Three-dimensional points along this strip are reconstructed using triangulation from the geometry of the two calibrated cameras.

The scanning system can acquire a 3D surface in approximately three minutes and typically produces around 300,000 3D points. The 3D points have an average spacing of 0.6 mm and an average depth accuracy of 0.5 mm over a typical surface area of  $200 \times 280$  mm. These detected 3D points are triangulated to obtain a 3D mesh that represents

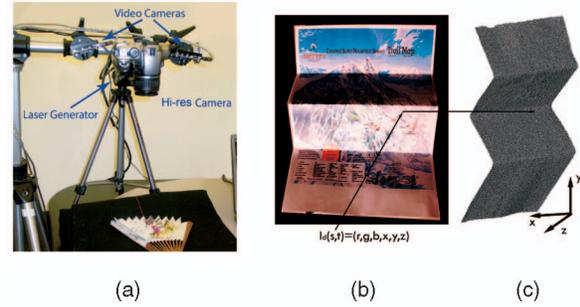


Fig. 3. (a) The 3D scanning system used to acquire the input data for our restoration framework. (b) The resulting depth image captured by our system. For each pixel  $I_d(s, t)$ , we have an associated RGB color pixel, as well as a 3D point  $(x, y, z)$ .

the document's surface. After the 3D surface is acquired, the high-resolution still camera is used to capture a 2D image of the document.

The triangulated 3D mesh can then be projected into the image of the still camera, which has a resolution of  $3,072 \times 2,048$ . Using the high-resolution image together with the reconstructed 3D points, we can create a *depth* image as follows:

$$I_d(s, t) = (r, g, b, x, y, z),$$

where  $(s, t)$  are coordinates in the depth image,  $(r, g, b)$  are the corresponding red-green-blue (RGB) pixel values, and  $(x, y, z)$  is the corresponding 3D values on the document's surface that is observed at position  $(s, t)$ . Fig. 3b shows an illustration of this depth image. Regions in the depth image that have no corresponding 3D points are ignored.

Storing the 3D and RGB data as one image makes it easy to reparameterize the 3D surface to a lower resolution if desired and serves to align the image gradients with geometry gradients in the illumination correction procedure.

### 4 GEOMETRIC CORRECTION

The depth map  $I_d(s, t)$  can be subsampled to produce a 3D mesh of desired resolution. This triangulated mesh is used in the conformal mapping procedure.

The basic idea of the conformal mapping procedure is that we want to map each triangle in the 3D mesh to a 2D plane while preserving angles. As such, under ideal conditions, our mapping would be a similarity mapping, allowing the 3D triangle to undergo change in translation, rotation, and scale. Therefore, we are solving for a set of 2D coordinates that minimizes the error in the piecewise similarity mappings for each triangle in the 3D mesh to a location in the 2D coordinate space while maintaining the mesh's topology. The mapping that globally minimizes this error will be taken to be the conformal map.

This differs from previous approaches that preserved vertex distances. In these previous algorithms, scale cannot change, as distances should be maintained. This results in a stronger constraint in the "flattening" procedure and leads to the need for iterative algorithms that compute the overall distance change per mesh edge at each step. Although the conformal mapping is a weaker constraint, we demonstrate

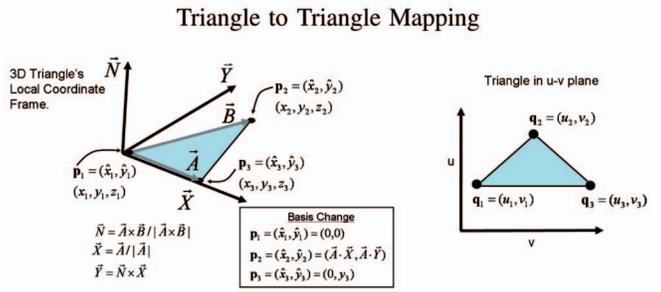


Fig. 4. Each triangle in the 3D mesh is converted from its full 3D representation  $(x_i, y_i, z_i)$  to its 2D representation in its local coordinate system defined by the triangles normal. The formula for conversion is given in the figure. We now consider the mapping of the 2D  $x - y$  local coordinates to the corresponding triangle in the  $uv$ -plane.

that it is just as effective for geometric correction and more efficient to obtain.

#### 4.1 Conformal Mapping

Consider a continuous closed 3D surface represented by a vector  $\phi(u, v)$ , parameterized by  $u$  and  $v$ , with components of  $x, y, z$ . The mapping  $R^3 \mapsto R^2$  of  $\phi(u, v)$  to the  $uv$ -plane is conformal [14] if and only if

$$\nabla^2 \phi = \frac{\partial^2 \phi}{\partial u^2} + \frac{\partial^2 \phi}{\partial v^2} = 0, \quad (1)$$

where  $\nabla^2$  is the Laplacian operator on  $\phi$ .

Since our input is a set of discrete 3D points reconstructed on the document's surface, we are looking for an inverse mapping from  $(x, y, z) \mapsto (u, v)$ . Consider now a 2D function  $f(x, y)$ , parameterized by  $(x, y)^1$  that returns a 2D point  $(u, v)$ . The 2D Laplacian operator can be satisfied by the Cauchy-Riemann equation [26] such that

$$\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} = 0,$$

which implies

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y} \quad \text{and} \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}. \quad (2)$$

To convert our problem to fit the form described in (2), the mapping of the 3D surface to a 2D coordinate frame can be reduced to  $R^2 \mapsto R^2$  by considering how to map the triangles of a 3D mesh to their corresponding triangles in the  $uv$ -plane [28]. Each 3D triangle is converted to its local 2D coordinate frame aligned by the triangle's normal as shown in Fig. 4. Using this idea, a mesh  $M = \{\mathbf{p}_{i \in [1, n]}, T_{j \in [1, m]}\}$  is generated, where  $\mathbf{p}_i$  are the  $n$  3D points acquired on the document's surfaces and serve as the mesh's vertices, and  $T_j$  represents the resulting  $m$  triangles, denoted as a tuple of three mesh vertices. The coordinates of the mesh vertices,  $\mathbf{p}_i$  can be expressed by their local coordinate frame, such that  $\mathbf{p}_i = (\hat{x}_i, \hat{y}_i)$ , which we will write for notation simplicity as  $\mathbf{p}_i = (x, y)$ .

A triangle-to-triangle mapping is defined by a unique affine transformation between the source and destination

triangle. If we consider the affine mapping  $f(\mathbf{p}) \mapsto \mathbf{q}$ , where  $\mathbf{p} = (x, y)$  and  $\mathbf{q} = (u, v)$ , where  $T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3}$  is the source triangle and  $T_{\mathbf{q}_1 \mathbf{q}_2 \mathbf{q}_3}$  is the destination, then the two triangles' vertices are related as

$$f(\mathbf{p}) = \frac{\text{area}(T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3}) \mathbf{q}_1 + \text{area}(T_{\mathbf{p}_2 \mathbf{p}_3 \mathbf{p}_1}) \mathbf{q}_2 + \text{area}(T_{\mathbf{p}_3 \mathbf{p}_1 \mathbf{p}_2}) \mathbf{q}_3}{\text{area}(T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3})}.$$

The partial derivatives (triangle gradient) of this equation is given as follows:

$$\begin{aligned} \frac{\partial f}{\partial x} &= \frac{\mathbf{q}_1(y_2 - y_3) + \mathbf{q}_2(y_3 - y_1) + \mathbf{q}_3(y_1 - y_2)}{2 \cdot \text{area}(T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3})}, \\ \frac{\partial f}{\partial y} &= \frac{\mathbf{q}_1(x_2 - x_3) + \mathbf{q}_2(x_3 - x_1) + \mathbf{q}_3(x_1 - x_2)}{2 \cdot \text{area}(T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3})}. \end{aligned} \quad (3)$$

This triangle gradient can be used to formulate the Cauchy-Riemann equations stated in (2). This can be written compactly in matrix form as follows:

$$\begin{bmatrix} \frac{\partial u}{\partial x} - \frac{\partial v}{\partial y} \\ \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y} \end{bmatrix} = \frac{1}{2A_T} \begin{bmatrix} \Delta x_1 & \Delta x_2 & \Delta x_3 & -\Delta y_1 & -\Delta y_2 & -\Delta y_3 \\ \Delta y_1 & \Delta y_2 & \Delta y_3 & \Delta x_1 & \Delta x_2 & \Delta x_3 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ \vdots \\ v_1 \\ v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

where  $\Delta x_1 = (x_3 - x_2)$ ,  $\Delta x_2 = (x_1 - x_3)$ ,  $\Delta x_3 = (x_2 - x_1)$ , and  $\Delta y_i$  is defined similarly from (3);  $A_T$  is the area of the triangle defined by  $T_{\mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3}$ . Solving this equation will find the appropriate  $(u_i, v_i)$  that are conformal on  $f(x, y) \mapsto (u, v)$ . Working from this equation, a global system of equations,  $\mathbf{A}\mathbf{x} = \mathbf{b}$ , can be written that incorporates all the vertices and triangles in mesh  $M$ , such that matrix  $\mathbf{A}$  and vector  $\mathbf{x}$  are defined as

$$\begin{bmatrix} a_{j,i} = \begin{cases} \frac{\Delta x_j}{2A_{T_j}}, & \text{if } \mathbf{p}_i \in T_j \\ 0, & \text{otherwise.} \end{cases} & \vdots & a_{j,2i} = \begin{cases} -\frac{\Delta y_j}{2A_{T_j}}, & \text{if } \mathbf{p}_i \in T_j \\ 0, & \text{otherwise.} \end{cases} \\ \dots & \dots & \dots \\ a_{2j,i} = \begin{cases} \frac{\Delta y_j}{2A_{T_j}}, & \text{if } \mathbf{p}_i \in T_j \\ 0, & \text{otherwise.} \end{cases} & \vdots & a_{2j,2i} = \begin{cases} \frac{\Delta x_j}{2A_{T_j}}, & \text{if } \mathbf{p}_i \in T_j \\ 0, & \text{otherwise.} \end{cases} \end{bmatrix} \begin{bmatrix} u_0 \\ \vdots \\ u_i \\ \vdots \\ u_n \\ \dots \\ v_0 \\ \vdots \\ v_i \\ \vdots \\ v_n \end{bmatrix} = \mathbf{0}.$$

The entries,  $a_{j,i}$  of matrix  $\mathbf{A}$ , can be described as follows: each triangle,  $T_j$ , occupies two matrix row entries located at row  $j$  and  $2j$ . For each triangle row pair, six entries per row will result from the vertices  $\mathbf{p}_k \in T_j$  occupying the columns  $k$  and  $2k$  with their corresponding  $\Delta x_k$  and  $\Delta y_k$ , as defined above. This results in  $\mathbf{A}$  being a  $2m \times 2n$  matrix, where  $m$  is the number of triangles, and  $n$  is the number of vertices in the 3D mesh. The entries for vector  $\mathbf{x}$  are the desired conformal points  $(u_i, v_i)$ .

1. Not to be mistaken with the surface  $(x, y, z)$  points.

### 4.2 Solving the System

To obtain a unique solution for the system  $Ax = b$  up to a similarity, some vertices must be constrained—otherwise, the  $(u_i, v_i)$  could have any arbitrary orientation in the 2D plane. Fixing at least two values will give a unique solution and constrains the orientation in the resulting conformal map. For example, if we want to constraint  $l$  vertices,  $p_{k'}$  to map to specified conformal locations  $q_{k'}$ , we modify the matrix  $A$  and vector  $b$  to reflect this constraint, such that  $A$ 's  $l$  columns entries,  $a_{k'}$  and  $a_{2k'}$ , are removed from the matrix  $A$ . A new  $b$  is constructed to incorporate this constraint in the solution, such that

$$b = - \begin{bmatrix} a_{k_1} & a_{2k_1} & \dots & a_{k_l} & a_{2k_l} \end{bmatrix} \begin{bmatrix} u_{k_1} \\ \vdots \\ u_{k_l} \\ v_{k_1} \\ \vdots \\ v_{k_l} \end{bmatrix},$$

where  $a_i$  are the columns removed from  $A$ , and  $(u_{k_i}$  and  $v_{k_i})$  represent the constrained conformal values.

After computing the new  $A$  and  $b$ , the linear system can be solved. Since, the solution is exact for only truly conformal surfaces, the vector  $x$  is computed to minimize  $\|Ax - b\|_2$ . This linear system can be solved using sparse solvers such as Conjugate Gradient methods. The resulting  $x$  vector contains the corresponding  $(u, v)$  values for the conformal parameterization.

### 4.3 Restoration Procedure

The location of each reconstructed 3D point,  $p_i$ , on the 3D mesh has an associated 2D location in the depth image,  $I_d(s_i, t_i)$ . After the conformal map has been constructed, we now have a mapping of the 3D mesh points to their conformal locations  $(u_i, v_i)$ . This implies that we also have a mapping from the distorted input image coordinates to the conformal parameterization, such that  $(s_i, t_i) \rightarrow (u_i, v_i)$ . The restoration procedure is now a matter of warping pixels in the input image to their locations in the restored image as shown in Fig. 5. Using the available topology of the 3D mesh

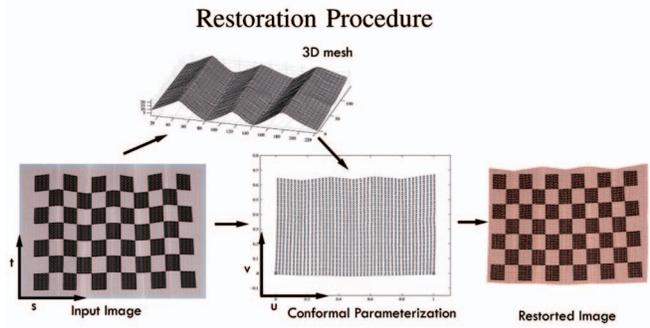


Fig. 5. Overview of the geometric correction procedure. Using the 3D depth image  $I_d(s, t)$ , a 3D mesh can be generated by sampling  $(s_i, t_i)$  points to obtain points on the document's surface. Conformal parameterization is performed to map the mesh to 2D coordinates. The corrected image is generated by warping the input  $(s_i, t_i)$  to their corresponding conformal coordinates  $(u_i, v_i)$ . This can be performed as simple texture mapping.

in the 2D conformal space (for example, see the mesh overlaid on the restore image in Fig. 5), this restoration procedure becomes a simple matter of texture mapping the distorted input image to the conformal space.

### 4.4 Specifying Fixed Vertices

As mentioned in Section 4.2, at least two of the vertices in the 3D mesh must be given fixed coordinates in the conformal space to ensure a unique solution. If no existing knowledge of the 3D mesh is available, then one reasonable choice is to fix the coordinates that correspond to the diameter of the mesh, that is, the two vertices in the mesh with farthest distance apart. Fig. 6 shows an example. The first example in the bottom row has no prior knowledge about the 3D point's relationship to image content. Arrows are drawn in Fig. 6 to denote the two points selected, these correspond to vertices that give the diameter of the mesh. These two corners are chosen to map to conformal locations  $[0, 0]$  and  $[1, 1]$ . In the second example, arrows are drawn to two vertices selected, vertices that correspond to the bottom corners of the mesh. Assume we know in advance that these points should lie on the lower horizontal edge in the output image, we can

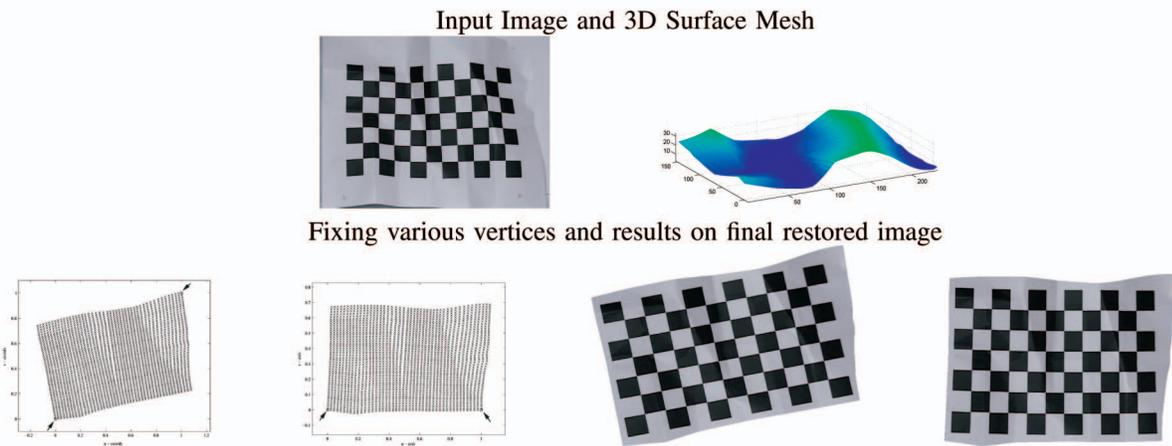


Fig. 6. Example of different constraints on the vertices. (Top) The distorted input image and its corresponding 3D mesh. (Bottom) The first image shows vertices corresponding to the diameter of the mesh being fixed to locations  $[0, 0]$  and  $[1, 1]$ . The second image shows the specification of two vertices that are known to line on the bottom horizontal axis of the document. These are fixed to be at locations  $[0, 0]$  and  $[1, 0]$ . The respective corrected images are shown next.

therefore constrain them to conformal positions  $[0, 0]$  and  $[1, 0]$ . This results in a horizontal orientation of the output conformal map. The two images based on these different constrained vertices are also shown in Fig. 6.

## 5 SHADING CORRECTION

Our shading correction uses the intrinsic image model which represents an image  $I$  as the product of the scene's reflectance  $R$  and the illumination  $L$ , that is,

$$I = L \cdot R. \quad (4)$$

Using this model, our scene reflectance,  $R$ , is assumed to be the 2D content printed on the document's surface. Recovering  $R$  will then allow us to produce a new  $I$  under uniform illumination. Although the intrinsic image model does not consider all illumination effects such as surface specularities, it is a reasonable model for our goals. As mentioned in Section 2, the intrinsic decomposition is ill posed given a single input image. The problem we try to address is to correct shading from a single image with known geometry but unknown arbitrary reflectance and illumination. It is, however, still an ill-posed problem. Even with known geometry, the lighting effect and surface reflectance are still interchangeable. Therefore, some further assumptions have to be made.

We assume that visually noticeable discontinuity in the imaged document, that is, an edge, is caused by either the reflectance of the printed content or by nonuniform illumination. Under this assumption, if we can identify and remove the image intensity change caused by the *illumination edges*, we can restore the original reflectance image (limitations of our assumption are discussed in Section 5.3). Thus, the problem of shading correction is reduced to a classification problem of illumination edges and reflectance (content) edges. Previous approaches have used machine learning techniques [4], [37] to perform this classification from a single image, but this problem is inherently ambiguous from only photometric information.

Our approach works from the 3D surface reconstruction and the captured RGB image of the document. Although the RGB tristimulus color space does have limitations in representing the true color properties of the imaged materials [23], this is currently the industry standard for commercial imaging hardware and our algorithm is therefore based on the RGB input. For computation on the image, such as computing image gradients, we use the RGB image values directly. However, when correcting the illumination, we make an assumption that the illumination change on the 3D surface is caused by a change in the pure white luminance in the scene and therefore only affects the luminance component of the input image. Thus, our RGB image is converted to the YUV color space, which decomposed the color into luminance (Y) and chrominance (UV) components and we perform our restoration procedure only on the Y channel. The processed Y channel is combined with the original UV's to generate the final image. It should be noted that this approach is not optimal since severe shading variations will result in changes in the U and V channels. Nevertheless, we found that this approach produces more visually appealing images than processing the RGB channels separately.

Since our framework has acquired the document's 3D geometry, we can identify illumination edges based on depth discontinuities on the document's surface. Recall that

our input depth image,  $I_d$  contains both RGB and depth information. For convention with other work, we use the variables  $(x, y)$  to denote coordinates in this image, as  $I_d(x, y)$ . Also, we refer to  $I(x, y)$  to mean the RGB components of the depth image only. We first define the notion of illumination edges and reflectance (content) edges:

- **Illumination edge.** An illumination edge is the union of connected *illumination pixels* that lie on a discontinuity in the RGB image *and* a discontinuity in the associated depth values. Simply stated, these are edges in the RGB image that correspond to edges (discontinuities) in the 3D geometry.
- **Reflectance edge.** A reflectance edge is the union of connected *reflectance pixels* that lie on a discontinuity in the RGB image only. Simply stated, these are edges in the RGB image that correspond to smooth 3D geometry.

Note that both types of edges correspond to visual edges in the RGB image.

### 5.1 Illumination Edge Detection

Given the depth image with its RGB values and corresponding per-pixel depth, we would like to detect illumination edges. By its definition, illumination edges are RGB image edges that correspond to discontinuous 3D geometry. We can find these geometric discontinuities by examining the *gradient*, denoted by the  $\nabla$  operator, of the RGB and depth values. Candidate illumination pixels can be detected by simple thresholding of the  $\nabla I_d$  depth image, computed using  $(x, y, z)$  values only, and their union with a thresholded  $\nabla I$  that considers  $(r, g, b)$  values only. Such operations can be robustly performed using existing gradient-based edge detectors such as a Canny edge detector. The result is a set of edges,  $\{E_i\}$ , correspond to illumination edges. For an example of a set extracted of illumination edges, see Fig. 10c.

### 5.2 Image Restoration

Our image restoration procedure starts with the image gradient field,  $\nabla I$ , of the input image. After illumination edge detection, we set the gradient of the illumination pixels (that is,  $\mathbf{p} \in \{E_i\}$ ) to zero to remove the illumination change at  $\mathbf{p}$  in  $\nabla I$ . We now wish to integrate the gradient field in order to reconstruct an image with illumination artifacts reduced. This can be achieved by solving the Poisson equation as demonstrated in [33], where a scalar function can be reconstructed from a guidance vector field and a given boundary condition.

Let  $\Omega$  stand for the image to be restored and  $w$  the guidance vector field (that is, gradient field). The Poisson equation with Dirichlet boundary conditions can be formulated as

$$\nabla^2 f = \nabla \cdot w, f|_{\partial\Omega} = f^*|_{\partial\Omega}, \quad (5)$$

where  $f$  is the unknown scalar function,  $f^*$  provides the desired values on the boundary  $\partial\Omega$ ,  $\nabla \cdot w = \frac{\partial w_x}{\partial x} + \frac{\partial w_y}{\partial y}$  is the divergence of  $w = (w_x, w_y)$ , and  $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$  is the Laplacian operator.

For our purpose, we only modify the gradient field, the original input image is used as the boundary condition. The idea behind our restoration can be considered as follows: The Poisson equation will produce a scalar field (that is, a new image) that tries to satisfy the boundary condition

while maintaining the specified gradient field. With our modified gradient field, pixels that lie on illumination edges should now produce a gradient field that is equal to zero, that is, the image region should no longer produce an edge. The pixel values in the input image around illumination edges will be modified to reduce the visible edge and therefore remove the gradient. At the same time, these changed pixels also try to maintain their initial values (the boundary condition). Consider now that this is performed globally over all input intensities and all gradient values, thus, pixel values in regions where the underlying gradient field has not been changed should not change, but those regions where the underlying gradient has been changed (illumination edges) should change. The result then is a new image similar to the input, where the illumination edges are reduced while maintaining the reflectance edges.

Our expected solution is therefore the unique solution of the discrete form of the Poisson equation (5), where  $f$  is discretized naturally using the underlying discrete pixel grid of the image,  $\Omega$ , and the vector guidance field  $w$  is set to be

$$w(\mathbf{p}) = \begin{cases} 0 & \text{if pixel } \mathbf{p} \in \{E_i\}, \\ \nabla I(\mathbf{p}) & \text{otherwise.} \end{cases} \quad (6)$$

In this form, the discrete Poisson equation is a sparse linear system that can be solved numerically [33]. Note that because we have zeroed out values in the gradient field that correspond to illumination edges, it is no longer guaranteed that the gradient field is integrable and, therefore, solving the Poisson equation only provides a solution in a least square sense in the recovered scalar field [2]. The consequence of this can be exhibited as a smoothing effect near illumination edges.

### 5.3 Bleed Correction

The techniques presented above make an assumption that the illumination edges and reflectance edges do not overlap. Although this assumption is made in other shading correction algorithms (for example, [27], [20]), it is rarely true in practice. When the two types of edges intersect or overlap, the restored images usually exhibits some color bleeding, that is, the reflectance (content) edges are blurry near shading boundaries; an example is shown in Fig. 7.

To ameliorate the bleeding problem, we compute its gradient direction in both the depth image and the RGB image. If the two directions differ by a large amount, we consider this to mean that there is an intersection of illumination and reflectance change. For such cases, we remove this pixel from the illumination edge pixel set by restoring its gradient value in the gradient field (6).

Fig. 7 demonstrates the effectiveness of this approach on a synthetic checkerboard pattern. The first image in Fig. 7 shows the input image. The second image shows the bleeding of pixel values where reflectance and illumination edges coincide. In this example, the gradient field at these pixel locations has been zeroed out because they were part of an illumination edge (shown as vertical lines). If we find that the direction of the RGB and depth gradients differ significantly (in our cases, we set the threshold for differences to be values greater than 30 degrees), the gradient field at this point is restored. This corresponds to the red points shown in the third image in Fig. 7. The results

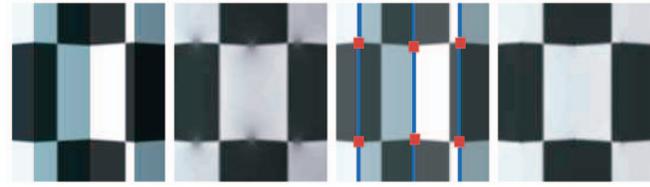


Fig. 7. Effects with and without bleed correction. From left to right: input image, restored image exhibiting bleeding, illumination mask with pixels that have been determined to be intersections of illumination and reflectance edges marked (as red points), and the result with bleed correction in place.

of making this adjustment are shown in the fourth image, which exhibits significantly less bleeding.

## 6 EXPERIMENTAL RESULTS

We have implemented our restoration framework outlined in Section 3 and have tested it on a variety of inputs. In this section, we first present quantitative results. In particular, we demonstrate our results compared with those of our previous work for geometric correction [6]. We show that we get virtually identical results to the distance-preserving approaches. Next, we demonstrate the shading correction on synthetic results to show improvements in terms of peak signal-to-noise ratios (PSNR). In the quantitative evaluation, geometric and photometric corrections are evaluated independently. These results are followed by qualitative results to demonstrate the results on items representative of a real-world digitization setup.

### 6.1 Quantitative Evaluation

**Geometric correction accuracy.** This example is intended to show the results as compared with distance-preserving flattening algorithms. In particular, we are using previous data presented in [6]. This data was prepared by first taking a control image before distortion and then distorting the document. The document's 3D information was captured using the structured-light scanner similar to that described in Section 3. The depth image is uniformly sampled to obtain  $46 \times 46$  3D points, resulting in a 3D mesh with 2,116 vertices and 4,050 triangles.

Fig. 8 shows five test cases. A document was imaged before it was warped (that is planar) and serves as the experiment's control. The five documents are restored using the conformal mapping procedure described in Section 4. To evaluate its geometric accuracy, we compare the distances between the *corners* of the checkerboard pattern in the restored and the controlled images. The corners are extracted automatically using a corner detector. There are 96 corners in total. A projective transform is used to align the restored document's and control's extracted corners. Our experiments show that the corner features in the restored images are within a pixel of the control patterns checkerboard corners.

We see that our results are virtually identical to those obtained in [6] using the distance-preserving approach. Although they have a slightly higher error in some test cases, of roughly 0.50 pixels, this should be taken into full context. The algorithms presented in [6] took over a minute to converge to a solution, whereas the conformal map takes a matter of seconds in an unoptimized Matlab code. The main purpose of this experiment is to show that the

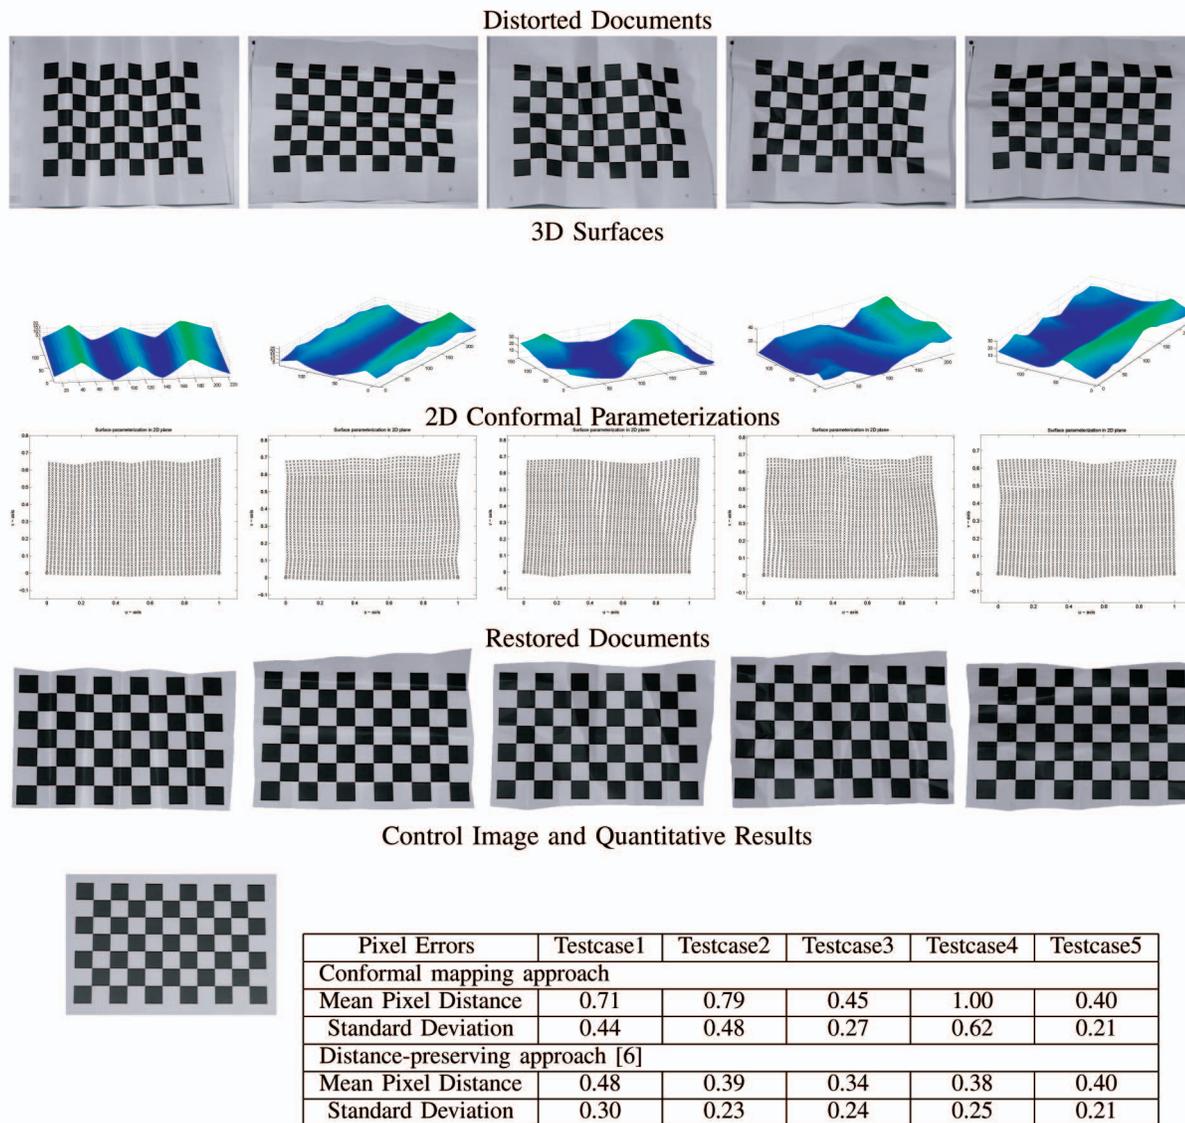


Fig. 8. Results of geometric correction algorithm applied to five test cases. The figure shows the imaged documents and their corresponding 3D meshes. The conformal mapping and resulting restored images are shown. The last row shows the experimental control and quantitative results. The mean pixel distance (and standard deviation) between the checkerboard corners of the *restored* and control images are given. This is compared with a distance-preserving approach.

conformal mapping approach can be successfully used in document restoration techniques that acquire a 3D representation of the document's surface. Thus, our approach will provide the same improvements to the performance in tasks such as OCR, as demonstrated in [6], [12].

**Photometric correction accuracy.** Synthetic images are used to evaluate the accuracy of photometric correction. Note that for these experiments, we do not correct the geometry, only the effects of nonuniform illumination. Fig. 9 shows five test cases, all with different distorted geometry and texture (that is, content), which are rendered in OpenGL with and without shading.

The image sets without shading serve as the experiment's control. The five documents with shading are then restored using the procedure described in Section 5. Treating the color deviation from the control image as "noise," the PSNR of the images before and after correction are computed. Although it is easy to see the perceptual

improvement of the correction, the experiments show that a quantitative improvement in PSNR (of up to several decibels) can also be obtained.

## 6.2 Qualitative Results

Fig. 10 shows an example of a distorted document that has been processed by our framework. The 2D content is of a checkerboard pattern. Fig. 10a shows the input image and 3D shape acquired by our system. Fig. 10b shows the resulting output image after geometric correction. In Fig. 10c, the top image shows the document's shading artifacts corrected and the bottom image shows a pixel mask denoted illumination edges. Fig. 10d shows the final output combining the results from both restoration algorithms. The document content looks significantly better than the input image. The lines are straight and the illumination appears uniform.

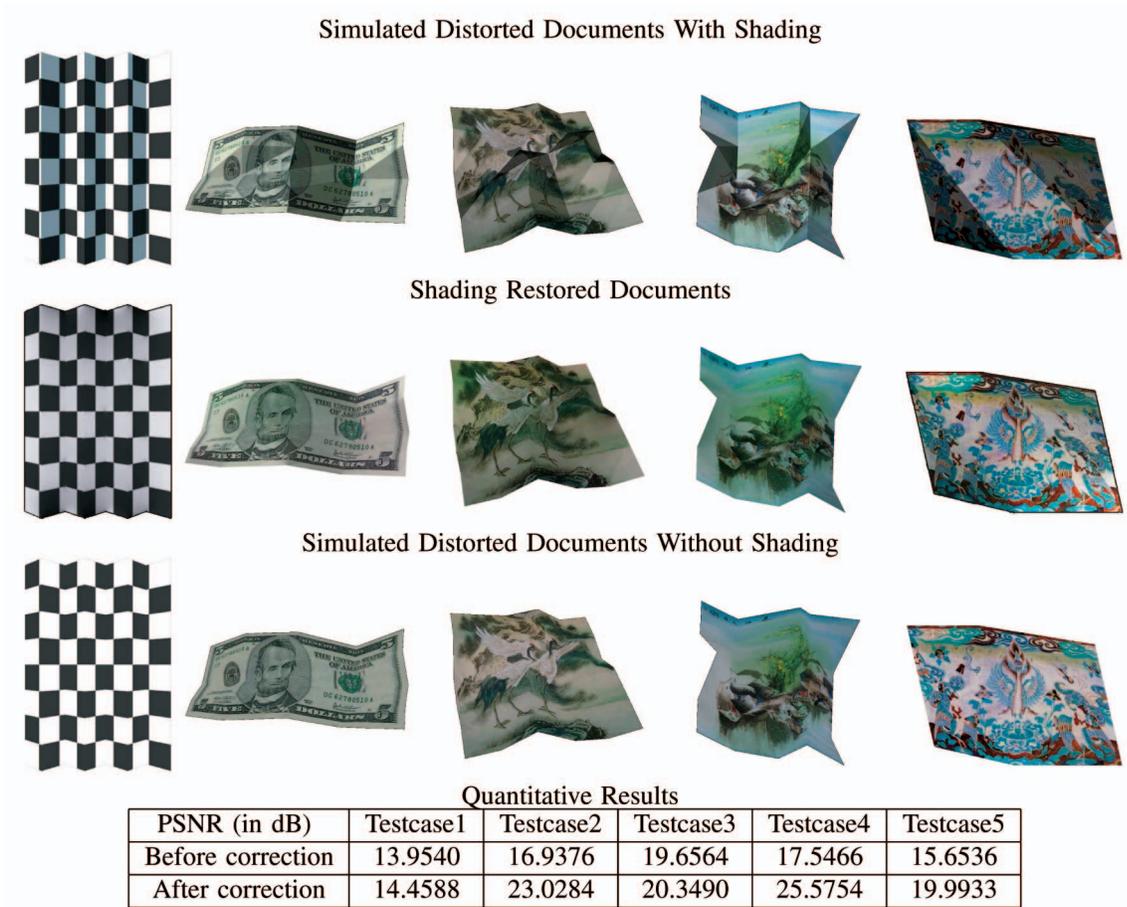


Fig. 9. Results of our photometric correction algorithm applied to five test cases. The input RGB images are decomposed to YUV, and the correction is applied on to the Y channel, which is combined with the original UV to obtain the output image. The figure shows the simulation images, generated in OpenGL, with/without shading in the first and third row separately, and the corresponding shading correction image are shown in the second row. The table shows quantitative results. The PSNR of the images before and after correction are given. PSNR is computed over the RGB space.

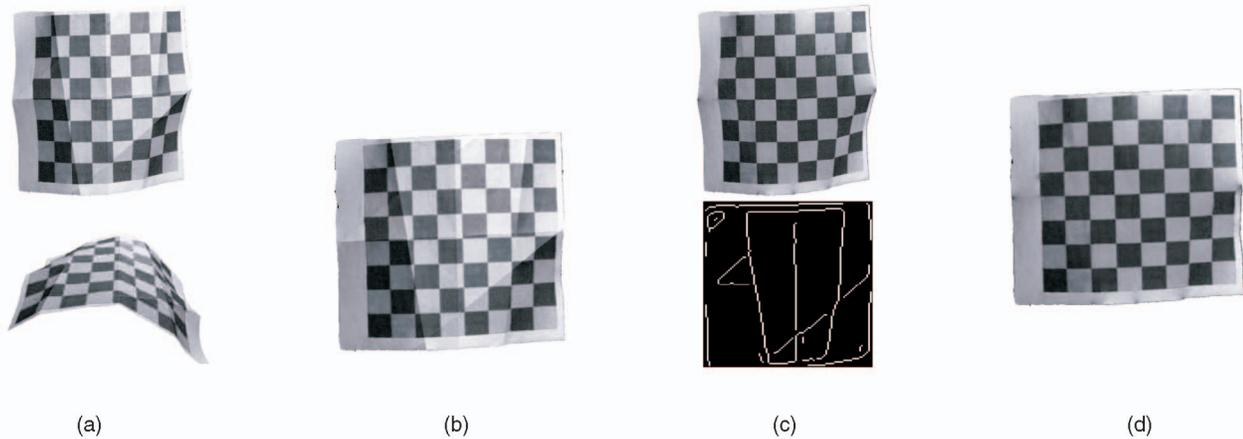


Fig. 10. Example of using our framework: (a) captured texture image and 3D geometry, (b) geometric distortion corrected, (c) input image with shading corrected and the illumination edge mask, and (d) results combined to generate the final output of the document without geometric and shading distortions.

Fig. 11 shows a *colored* example. This is an example of a map that has been folded at one time and is not completely unfolded when imaged. This example has significant shading artifacts. The first image shows the input image. The second shows the acquired 3D mesh. The third image shows the geometry rectified with the original input shading. The last

image shows the final recovered 2D content. In this example, simple thresholding or image filter is not applicable to correct the shading. Our additional shading correction provides a significant improvement in the restoration.

Our next example is an oriental fan shown in Fig. 12. The design of the fan itself is problematic for imaging. Fig. 12



Fig. 11. (From left to right): original image, 3D mesh, geometry corrected, and geometry and illumination corrected.

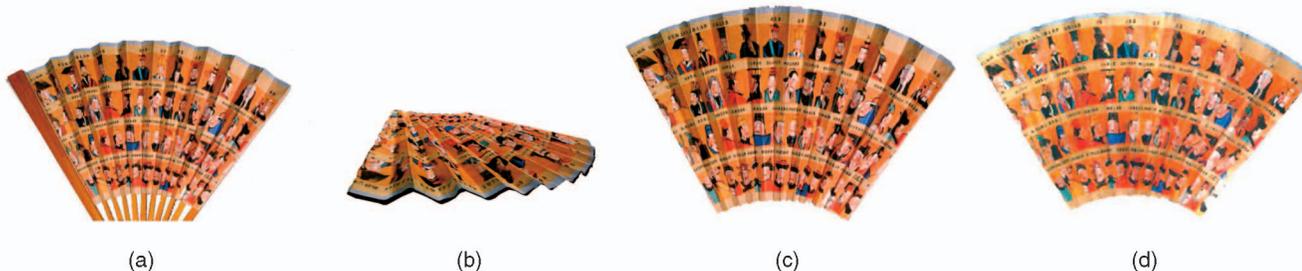


Fig. 12. An oriental folding fan. (a) Input original image, (b) 3D view of the scanned data, (c) fan geometry rectified, and (d) final image with geometry and shading corrected image.

shows the original input, the 3D surface reconstruction, the results from geometric rectification and the results from both geometric and shading rectification. Although some bleeding is evident in the shading correction, the final image is nonetheless an improvement over the original input.

## 7 DISCUSSION

Our system is useful in correcting both geometric and shading distortion. The nature of our approach does assume that the cameras can observe the entire surface, that is, there are no self-occlusions on the document surface. In addition, accuracy in the geometry distortion is related to sampling of the 3D mesh. In the experiments with the five checkerboard patterns in Section 6, we use a simple mesh of  $46 \times 46$  vertices, which was sampled uniformly over the depth image. In practice, much better (and higher) sampling that better preserves the documents surface can be used. Having said this, we also find in practice that even low-resolution sampling can give sufficient results. For example, experiments presented in [34] used only  $20 \times 15$  surface samples to reconstruct the document's surface with acceptable results.

Although the algorithm can correct geometric distortion, it cannot fill in missing intensity information lost due to projection. Therefore, regions of high deformation may appear blurry in the restored image due to a lack of intensity information. To address this problem, multiple images of the distorted document would need to be captured and registered to fill in missing intensity information.

Such geometric and image sampling issues are problematic for all restoration approaches based on geometric models. Instead, our goal is to show that the conformal parameterizations of the document's 3D surface can be used to restore arbitrarily distorted documents to within a *single* pixel of their true planar representation. These results are virtually identical to previous techniques based on distance-preserving algorithms; however, the conformal map can be computed in a matter of seconds.

With respect to correcting shading, our formulation cannot deal with smooth shading changes or those caused by shadows. Smooth shading changes fit well with the *Retinex* theory [27] and can therefore be removed by many existing techniques such as the *bilateral filtering* technique [32], which is relatively simple to implement. Sharp shadow edges are more challenging to remove. Under certain lighting assumptions, it has been shown that hard shadows can be successfully removed by first finding a shadow invariant image [17]. Although we have not implemented these techniques, we envision that the resulting image from our approach can be further processed to remove these shading artifacts we currently do not address.

Although 3D scanning may seem excessive for document imaging, for rare and valuable documents, for example, those found in special collections and museums, such processing is often warranted. In addition, these types of documents are more likely to suffer from geometric distortion and require special imaging approaches. Very old documents, which are often brittle and dry, are at risk for tearing and breaking when physically handled. Our system provides a viable alternative to physical restoration for such items.

## 8 CONCLUSION

We have demonstrated a system for correcting geometric and photometric distortion in the image of arbitrarily distorted documents. Our approach acquires a 3D reconstruction of the document together with a registered high-resolution 2D image. Conformal mapping is used to remove geometric distortion by mapping the 3D surface to a plane while minimizing angular distortion. Illumination correction is performed by using the 3D shape in order to distinguish between *content* gradient edges and *illumination* gradient edges. The results of these two processes are combined to

remove geometric and photometric artifacts from images of distorted documents. As demonstrated in this paper, our approach can significantly improve the visual appearance of the content printed on distorted documents.

## REFERENCES

- [1] G. Agam and C. Wu, "Structural Rectification of Non-Planar Document Images: Application to Graphics Recognition," *Lecture Notes in Computer Science*, p. 2390, 2002.
- [2] A. Agrawal, R. Raskar, and R. Chellappa, "What Is the Range of Surface Reconstructions from a Gradient Field," *Proc. European Conf. Computer Vision*, 2006.
- [3] H.G. Barrow and J.M. Tenenbaum, *Recovering Intrinsic Scene Characteristics from Images*. Academic Press, 1978.
- [4] M. Bell and W.T. Freeman, "Learning Local Evidence for Shading and Reflection," *Proc. Int'l Conf. Computer Vision*, pp. 670-677, 2001.
- [5] A. Blake, "Boundary Conditions for Lightness Computation in Mondrian World," *Computer Vision, Graphics and Image Processing*, vol. 32, pp. 314-327, 1985.
- [6] M.S. Brown and W.B. Seales, "Document Restoration Using 3D Shape: A General Deskewing Algorithm," *Proc. Int'l Conf. Computer Vision*, July 2001.
- [7] M.S. Brown and W.B. Seales, "Image Restoration of Arbitrarily Warped Documents," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 10, pp. 1295-1306, Oct. 2004.
- [8] M.S. Brown and Y.C. Tsoi, "Geometric and Shading Correction for Images of Printed Materials Using Boundary," *IEEE Trans. Image Processing*, vol. 15, no. 6, pp. 1544-1554, 2006.
- [9] M.S. Brown and C.J. Pisula, "Conformal Deskewing of Warped Documents," *Computer Vision and Pattern Recognition*, 2005.
- [10] H. Cao, X. Ding, and C. Liu, "A Cylindrical Surface Model to Rectify the Bound Document Image," *Proc. Int'l Conf. Computer Vision*, pp. 228-233, 2003.
- [11] H. Cao, X. Ding, and C. Liu, "Rectifying the Bound Document Image Captured by the Camera: A Model Based Approach," *Proc. Int'l Conf. Document Analysis and Recognition*, 2003.
- [12] K.B. Chua, L. Zhang, Y. Zhang, and C.L. Tan, "A Fast and Stable Approach for Restoration of Warped Document Images," *Proc. Int'l Conf. Document Analysis and Recognition*, 2005.
- [13] U. Clarenz, N. Litke, and M. Rumpf, "Axioms and Variational Problems in Surface Parameterization," *Computer Aided Geometric Design*, vol. 21, no. 8, pp. 727-749, 2004.
- [14] R. Courant, "Plateau's Problem," *Dirichlet's Principle, Conformal Mapping and Minimal Surfaces*, chapter 3, pp. 95-134, Springer, 1977.
- [15] P. Debevec, "Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography," *Proc. Int'l Conf. Computer Graphics and Interactive Techniques*, pp. 189-198, 1998.
- [16] M. Desburn, M. Meyer, and P. Alliez, "Intrinsic Parameterizations of Surface Meshes," *Proc. Ann. Conf. European Assoc. for Computer Graphics*, vol. 21, pp. 209-218, Sept. 2002.
- [17] G. Finlayson, S. Hordley, and M. Drew, "Removing Shadows from Images," *Proc. European Conf. Computer Vision*, pp. 823-836, 2002.
- [18] M. Floater and M. Reimers, "Meshless Parameterization and Surface Reconstruction," *Computer Aided Design*, vol. 18, no. 2, pp. 77-92, Mar. 2001.
- [19] A. Fournier, A. Gunawan, and C. Romanzin, "Common Illumination between Real and Computer Generated Scenes," *Proc. Graphics Interface*, 1993.
- [20] B.V. Funt, M.S. Drew, and M. Brockington, "Recovering Shading from Color Images," *Proc. European Conf. Computer Vision*, pp. 123-132, 1992.
- [21] N. Gumerov, A. Zandifar, R. Duraiswami, and L.S. Davis, "Structure of Applicable Surfaces from Single Views," *Proc. European Conf. Computer Vision*, 2004.
- [22] B.K.P. Horn, "Determining Lightness from an Image," *Computer Vision, Graphics and Image Processing*, vol. 3, pp. 277-299, 1974.
- [23] G.M. Johnson and M.D. Fairchild, "Full-Spectral Color Calculations in Realistic Image Synthesis," *Computer Graphics and Applications*, vol. 19, no. 4, pp. 47-53, July/Aug. 1999.
- [24] T. Kanungo, "Document Degradation Models and Methodology for Degradation Model Validation," PhD dissertation, Dept. of Electrical Eng., Univ. of Washington, Mar. 1996.
- [25] T. Kanungo, R.M. Haralick, and I. Phillips, "Global and Local Document Degradation Models," *Proc. Int'l Conf. Document Analysis and Recognition*, pp. 730-734, 1993.
- [26] S.C. Krantz, "The Cauchy-Riemann Equations," *Handbook of Complex Analysis*, Birkhauser, 1999.
- [27] E.H. Land and J.J. McCann, "Lightness and Retinex Theory," *J. Optical Soc. Am.*, vol. 61, pp. 1-11, 1971.
- [28] B. Lévy, S. Petitjean, R. Nicholas, and J. Maillot, "Least Squares Conformal Maps for Automatic Texture Atlas Generation," *Proc. Int'l Conf. Computer Graphics and Interactive Techniques*, pp. 362-371, 2002.
- [29] J. Liang, D. DeMenthon, and D. Doermann, "Flattening Curved Documents in Images," *Proc. Computer Vision and Pattern Recognition*, pp. 338-345, 2005.
- [30] J. Liang, D. Doermann, and H. Li, "Camera-Based Analysis of Text and Documents: A Survey," *Int'l J. Document Analysis and Recognition*, vol. 7, no. 2-3, pp. 84-104, July 2005.
- [31] S. Marschner and D. Greenberg, "Inverse Lighting for Photography: 74217," *Proc. IS&T/SID Fifth Color Imaging Conf.*, pp. 262-265, 1997.
- [32] B.M. Oh, M. Chen, J. Dorsey, and F. Durand, "Image-Based Modeling and Photo Editing," *Proc. Int'l Conf. Computer Graphics and Interactive Techniques*, pp. 433-442, 2001.
- [33] P. Perez, M. Gangnet, and A. Blake, "Poisson Image Editing," *Proc. Int'l Conf. Computer Graphics and Interactive*, vol. 22, no. 3, pp. 313-318, 2003.
- [34] M. Pitu, "Undoing Paper Curl Distortion Using Applicable Surfaces," *Proc. Computer Vision and Pattern Recognition*, Dec. 2001.
- [35] M.X. Sun, R.G. Yang, Y. Lin, G. Landon, B. Seales, and M.S. Brown, "Geometric and Photometric Restoration of Distorted Documents," *Proc. Int'l Conf. Computer Vision*, 2005.
- [36] C.L. Tan, L. Zhang, Z. Zhang, and T. Xia, "Restoring Warped Document Images through 3D Shape Modeling," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 2, pp. 195-208, Feb. 2006.
- [37] M.F. Tappen, W.T. Freeman, and E.H. Adelson, "Recovering Intrinsic Images from a Single Image," *Advances in Neural Information Processing Systems*, 2003.
- [38] R. Tsai, "An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision," *Proc. Computer Vision and Pattern Recognition*, pp. 364-374, 1986.
- [39] Y.C. Tsoi and M.S. Brown, "Geometric and Shading Correction of Imaged Printed Materials: A Unified Approach Using Boundary," *Proc. Computer Vision and Pattern Recognition*, June 2004.
- [40] A. Ulges, C.H. Lampert, and T.M. Breuel, "Document Capture Using Stereo Vision," *Proc. ACM Symp. Document Eng.*, pp. 198-200, 2004.
- [41] A. Ulges, C.H. Lampert, and T.M. Breuel, "Document Image Dewarping Using Robust Estimation of Curled Text Lines," *Proc. Int'l Conf. Document Analysis and Recognition*, 2005.
- [42] T. Wada, H. Ukida, and T. Matsuyama, "Shape from Shading with Interreflections under Proximal Light Source," *Proc. Int'l Conf. Computer Vision*, pp. 66-71, 1995.
- [43] Y. Weiss, "Deriving Intrinsic Images from Image Sequences," *Proc. Int'l Conf. Computer Vision*, vol. 2, pp. 68-75, July 2001.
- [44] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," *Proc. Int'l Conf. Computer Graphics and Interactive Technique*, pp. 215-224, 1999.
- [45] L. Zhang, Z. Zhang, C.L. Tan, and X. Tao, "3D Geometric and Optical Modeling of Warped Documents Images from Scanners," *Proc. Computer Vision and Pattern Recognition*, 2005.
- [46] Z. Zhang and C.L. Tan, "Correcting Document Image Warping Based on Regression of Curved Text Lines," *Proc. Int'l Conf. Document Analysis and Recognition*, 2003.
- [47] Z. Zhang, C.L. Tan, and L. Fan, "Estimation of 3D Shape of Warped Documents Surface for Image Restoration," *Proc. Computer Vision and Pattern Recognition*, 2004.



**Michael S. Brown** received the BS and PhD degrees from the University of Kentucky in 1995 and 2001, respectively. He was a visiting PhD student at the University of North Carolina—Chapel Hill in 1998-2000. He is currently an assistant professor at Nanyang Technological University in Singapore and has held the position of assistant professor at Hong Kong University of Science and Technology and California State University—Monterey Bay. He is a member of the IEEE.



**Yun Lin** received the BS and MS degrees in computer science from HuaZhong Normal University, China. Currently, she is a PhD candidate in computer science at the University of Kentucky. Her research interests include image processing and computer graphics and visualization with applications in the area of digital libraries and medical images. She is a student member of the IEEE.



**Mingxuan Sun** received the BS degree from Zhejiang University, China, in 2004 and the MS degree in computer science from the University of Kentucky in 2006. Currently, she is a PhD student in computer science at the Georgia Institute of Technology. Her research interests include shading correction and light field rendering. Prior to Georgia Tech, she has worked in the Center for Visualization and Virtual Environments at the University of Kentucky.



**W. Brent Seales** received the PhD degree in computer science from the University of Wisconsin—Madison. He is an associate professor of computer science and the research director at the Center for Visualization and Virtual Environments at the University of Kentucky. His research interests include in computer vision and image processing, with applications in digital libraries, surgical technology and medical visualization, and multimedia. He is a member of the IEEE.



**Ruigang Yang** received the PhD degree in computer science from the University of North Carolina—Chapel Hill in 2003 and the MS degree in computer science from Columbia University in 1998. He is an assistant professor in the Computer Science Department, University of Kentucky. His research interests include computer graphics, computer vision, and multimedia. He is a recipient of the US National Science Foundation (NSF) CAREER award in 2004 and a member of the IEEE and ACM.

▷ For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).