Geometric and Shading Correction for Images of Printed Materials A Unified Approach Using Boundary

Yau-Chat Tsoi and Michael S. Brown
Department of Computer Science
H.K.U.S.T.
Clear Water Bay, Hong Kong
{desmond, brown}@cs.ust.hk

Abstract

We present a novel approach that uses boundary interpolation to correct (1) geometric distortion and (2) shading artifacts present in images of printed materials. Unlike existing approaches, our algorithm can simultaneously correct a variety of geometric distortions, including skew, fold distortion, binder curl, and combinations of these. In addition, the same interpolation framework can estimate the *intrinsic* illumination component of the distorted image to correct shading artifacts.

1. Background and Contribution

Printed materials are digitized for electronic dissemination and archival. While the printed content is 2D, the physical media that the content is printed on is rarely 2D. As a result, two types of distortion are commonly present in the images of these materials. The first is geometric distortion of the 2D content arising from material's non-planar shape. The second is shading artifacts also resulting from the non-planar shape. Examples of these effects can be seen in Figure 1.

Geometric and shading distortion is particularly troublesome for images of large materials, such as oversize books and art-like materials. Such items are unsuitable for flatbed imaging and are instead imaged using high resolution cameras. As a result, they are rarely flattened before imaging. Moreover, some materials simply cannot be made completely flat without risk of damage. This makes it difficult to avoid geometric distortion and shading artifacts. For these materials, distortion correction must be applied after imaging.

Previous work has focused primarily on two types of distortions: *skew* and *binder-curl*. Skew arises when the acquired image and the physical material's 2D content are not axis aligned. Approaches [1, 9, 10, 12] to correct this distortion compute planar transformations to rectify the imaged content. Binder-curl effect occurs in the region near a book's spine. Binder-curl results in an unsightly *curling* appearance in the 2D content. Approaches to undo this distortion [4, 5, 11, 15, 17] parameterize the distortion using cylindrical models which can be used to "unroll" the curved imagery.

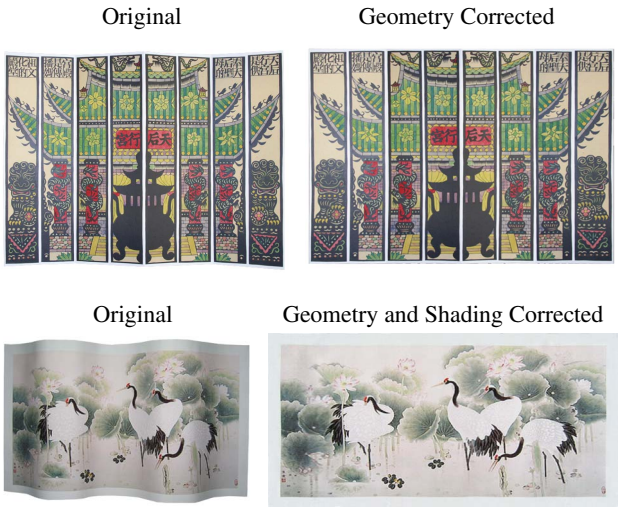

Figure 1: (Top) An image of a large piece of art that suffers from fold distortion before and after being corrected with our algorithm. (Bottom) An image of a painting suffering for geometric distortion and shading artifacts before and after correction.

These existing approaches are limited to planar documents or a single page of a bound volume. These approaches also assume the presence of only a single type of distortion. As a result, the skew and binder-curl correction are typically applied in tandem to remove all distortion. Furthermore, existing approaches assume that the targeted distortion is exhibited only once in the image. De-skewing algorithms cannot correct a material that exhibits multiple skewing (for example from folds), such as shown in Figure 1 (Top). Binder-curl approaches assume the image is of a single page and the curl is either on the left or right of the image. The image in Figure 1 (Bottom) with several curled regions cannot be corrected. None of the existing approaches can work for materials that exhibit both multiple skew (folds) and binder-curl.

0-7695-2158-4/04 $20.00 (C) 2004 IEEE

Another drawback to existing approaches is they do not address shading artifacts. Since these techniques target images of text-based documents and bound books, the acquired images are subsequently converted to bi-tonal representations for use in further processing, such as optical character recognition (OCR). Thus, shading artifacts are only considered in the binarization process, requiring the use of local vs. global thresholding to overcome variations in shading [17].

Recent approaches have addressed distortions from *arbitrarily* shaped materials [3, 13]. These approaches acquire a 3D reconstruction of the imaged material's surface in addition to a 2D image. Relaxation algorithms are used to flatten the 3D surface to a plane while minimizing spring energies between acquired 3D points. While these approaches work, the additional 3D shape information is rarely available and requires modifications to the imaging technology. In addition, these relaxation approaches are computationally slow [3, 13] and do not address shading artifacts.

[Our Contribution] In this paper, we present a unified approach to correct both geometric and photometric (shading) distortion via the 2D boundary of the imaged material. Using boundary interpolation, we can compute a corrective mapping to simultaneously undo common geometric distortions, such as skew, binder curl, and fold distortion (and combinations of these). In addition, the same interpolant framework can be used to estimate the intrinsic illumination image. This estimated illumination image together with the original image can be used to remove shading artifacts. Our proposed approach is fast, works directly from a 2D image, and provides a general solution to correct both simple and complex geometric distortions and shading.

The remainder of this paper is as follows: Section 2 and 3 overview our approach; Section 4 shows results using several examples; and Section 5 concludes our work.

2. Geometric Distortion Correction

2.1. Distortion Parameterization By Boundary

Our work assumes that the printed content will be rectilinear in its true planar format. Our goal is to find a parameterization between the desired rectilinear representation and the distorted input image. To do this, we consider how to model the material's 3D structure and how this model behaves under projection (imaging).

Several authors [4, 5, 15] have proposed cylindrical models to model binder-curl distortion of a single page. We use a more general model of a ruled surface composed of two *opposite-boundary* curves [7]. To visualize this model, consider a book distorted by binder curl represented in Figure 2. Given the top and bottom 3D boundary curves, $\mathbf{C}_1(u)$ and $\mathbf{C}_2(u)$, the entire 3D surface can be described as:

$$\mathbf{S}(u,v) = (1-v)\mathbf{C}_1(u) + v\mathbf{C}_2(u) \quad (1)$$

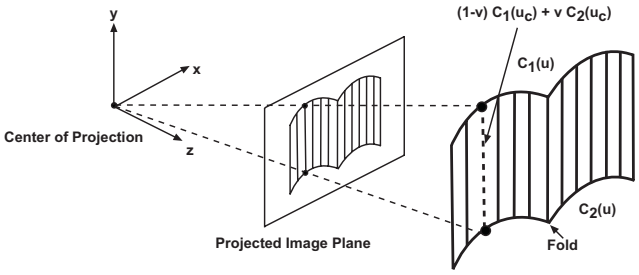

Figure 2: Material represented by a ruled surface defined by two opposite curves, $\mathbf{C}_1(u)$ and $\mathbf{C}_2(u)$. The projection of this surface can also be parameterized by the resulting 2D boundary curves.

where $\mathbf{S}(u,v)$ is the equation for the 3D ruled surface, parameterized by u and v. Given $\mathbf{S}(u,v)$, every iso-parametric line $u = u_c$ is a straight line segment between $\mathbf{C}_1(u)$ and $\mathbf{C}_2(u)$ parameterized by v. Note, in the case of folds, if the fold is along the iso-parametric lines $u = u_c$, the surface can still be modeled using only the boundary curves. Not only does this model describe the 3D structure, it also fits our notion of how printed materials behavior. Ruled surfaces are developable surfaces and can be mapped to a plane without distortion. This is similar to our notion of a printed page; it can be flattened without introducing any distortion to the 2D content.

While restrictive, this two-boundary ruled surface model is sufficient to model the *vast majority* of printed materials. Existing cylindrical models can be represented with our model. In the case of folded materials, the folds made are almost exclusively along iso-parametric lines to allow the material to be folded into equal pieces for storage. Thus, a document with several folds and curl effects can be modeled. We note that it is easy to violate this surface model; e.g. folding a page's corner over on itself introduces a fold against the iso-parametric lines, however such examples are not typical of the majority of imaged materials.

With the surface model in place, we examine its behavior under projection. The 2D projection of points on the two 3D boundary curves can be expressed as, $\mathbf{x}_1 = \tilde{\mathbf{P}}[\mathbf{C}_1(u_c)\ 1]^T$ and $\mathbf{x}_2 = \tilde{\mathbf{P}}[\mathbf{C}_2(u_c)1]^T$, where $\tilde{\mathbf{P}}$ is the 3×4 projection matrix of the camera. Since lines are preserved under projection, the line segments passing through, $\mathbf{C}_1(u_c)$ and $\mathbf{C}_2(u_c)$ must pass through the 2D boundary points $\mathbf{x}_1$ and $\mathbf{x}_2$ as shown in Figure 2. Thus, the projection of this model can also be parameterized using the projected boundary curves in 2D.

While only the two opposite boundary curves are needed to parameterize the distortion, in practice we will use four boundary curves corresponding to the rectilinear edges of the imaged materials. With appropriate boundary interpolation, it is not necessary to explicitly denote which curves are the opposite boundary pair. In addition, small nonlinearities in the imaging system, such as radial distortion, can be compensated for in the interpolation process.

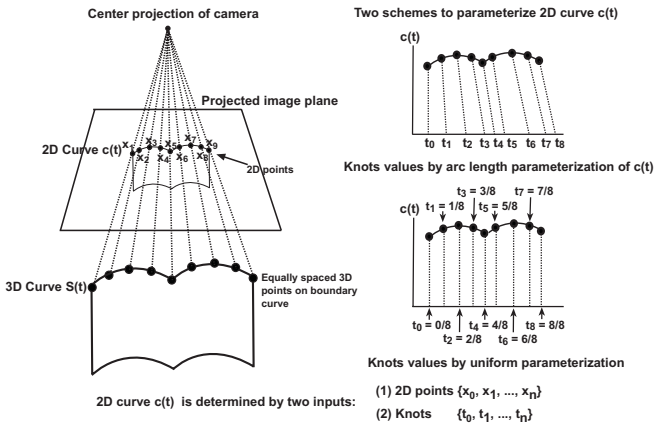

Figure 3: The 2D curve $\mathbf{c}(t)$ is defined by $n+1$ 2D feature points $\mathbf{x}_i$ *and* $n+1$ knot values t_i. Two schemes are used to compute the knot's values. **Arc length** parameterization assigns knot values using their position along the arc length of $\mathbf{c}(t)$. **Uniform** parameterization assumes the knots are uniformly sampled in parameter space, i.e. $t_i = i/n$. If $\mathbf{x}_i$ are from equally sampled 3D points along the 3D curve, uniform parameterization can encode useful depth information.

2.2. Boundary Representation

Four boundary curves of the distorted material are represented by curves $\mathbf{c}_1$, $\mathbf{c}_2$, $\mathbf{c}_3$, $\mathbf{c}_4$, representing the top, right, bottom and left sides of the material. Each of these curves is encoded using a natural cubic spline that defines continuous image coordinates along the curve. i.e. $\mathbf{c}_i = (x(t), y(t))$, where t is a parameter along the curve with range between 0 and 1.

Recall that 2D natural cubic splines (NCS) are defined by specifying a set of $n+1$ 2D points, $\mathbf{x}_0$, $\mathbf{x}_1$, ..., $\mathbf{x}_n$ together with their corresponding parameter value (knots), t_0, t_1, ..., t_n. The 2D points $\mathbf{x}_i$ and knots t_i, uniquely define a set of piece-wise functions, $\mathbf{s}_i(i)$ such that:

$$\mathbf{s}(t) = \begin{cases} \mathbf{s}_0(t) & t \in [t_0, t_1] \\ \mathbf{s}_1(t) & t \in [t_1, t_2] \\ \vdots & \\ \mathbf{s}_{n-1}(t) & t \in [t_{n-1}, t_n] \end{cases}, \tag{2}$$

where $\mathbf{s}_i(t) = (x_i(t), y_i(t))$, and $x_i(t)$ and $y_i(t)$ are cubic functions. The coefficients of these cubic pieces can be uniquely computed with the constraints that $\mathbf{s}_{i-1} = \mathbf{x}_i = \mathbf{s}_i$ for $(1 \le i \le n-1)$ and that $\mathbf{s}'$ and $\mathbf{s}''$ are continuous.

While the 2D points $\mathbf{x}_0$, $\mathbf{x}_1$, ..., $\mathbf{x}_n$ can be obtained by extracting image coordinates along the boundary, a corresponding set of knot values, t_i must also be specified. Two different parameterization approaches are used.

The first approach uses the approximated 2D arc-length, $|\mathbf{s}|$ of $\mathbf{s}(t)$ which is defined by the chord length of the points $\mathbf{x}_i$, such that each t_i is defined as:

$$|\mathbf{s}| = \sum_{0<i<n-1} |\mathbf{x}_i - \mathbf{x}_{i+1}|_2$$
$$t_i = \begin{cases} 0 & \text{if } i = 0 \\ \frac{1}{|\mathbf{s}|} \sum_{0<j<i} |\mathbf{x}_{j-1} - \mathbf{x}_j|_2 & \text{if } i > 0, \end{cases} \tag{3}$$

where $|\cdot|_2$ represents Euclidean distance. The idea of this parameterization is to move along the curve $\mathbf{s}(t)$ at a constant rate.

The second approach defines the parameterization *uniformly* with values $t_i = i/n$, where $n+1$ is the number of control points. In this case, uniform parameterization encodes how to move along the 2D curve as it corresponds to a uniform sampling along the *arc-length* of the actual 3D curve. Consider a 3D curve called $\mathbf{S}(t)$, where t is parameterized by the 3D curve's *arc-length*. If we have the relationship:

$$\mathbf{x}_i = \tilde{\mathbf{P}}[\mathbf{S}(t = i/n) \;\; 1]^T, \tag{4}$$

where $\tilde{\mathbf{P}}$ is the projection matrix of the camera, i is the i^{th} sample, and $n+1$ is the number of samples, then $\mathbf{x}_i$ represents the projection of the 3D curves points at $t_i = i/n$ in $\mathbf{S}(t)$. This can be visualized in Figure 3. The features points $\mathbf{x}_i$ are obtained from the uniform *sampling* of the 3D curve and corresponds to a uniform *parameterization* of the projected 2D curve. Such parameterization can be used if we have prior knowledge that 2D image points are from 3D points samples by arc-length in 3D. This 2.5D information incorporates the depth change information of the 3D curve into the parameterization.

2.3. Boundary Interpolation

The restored image is defined over the parametric space u and v, where $u \in [0,1]$ and $v \in [0,1]$. Each curve $\mathbf{c}_i$ maps to its corresponding side of the rectilinear image (shown in Figure 4). For example, the (x, y)'s points along the top boundary $\mathbf{c}_1(u)$ should map to $I(u, 0)$, in the undistorted image. This curve-to-line mapping can be computed by simple sampling of 2D curve.

While the boundaries can be easily mapped to their correct locations, a 2D function to describe how to map (x, y) points inside the $\mathbf{c}_i$ curves is needed. This 2D function can be provided using a bi-linearly blended Coons patch [6], as follows:

$$\begin{aligned} \mathbf{c}(u,v) = &[1-u \quad u] \begin{bmatrix} \mathbf{c}_4(v) \\ \mathbf{c}_2(v) \end{bmatrix} \\ &+ [\mathbf{c}_1(u) \quad \mathbf{c}_3(u)] \begin{bmatrix} 1-v \\ v \end{bmatrix} \\ &- [1-u \quad u] \begin{bmatrix} \mathbf{c}_1(0) & \mathbf{c}_2(0) \\ \mathbf{c}_3(1) & \mathbf{c}_4(1) \end{bmatrix} \begin{bmatrix} 1-v \\ v \end{bmatrix} \end{aligned} \tag{5}$$

Equation 5 is formed by a linear interpolation of two opposite-boundary curves (first two terms), with a corrective function based on the boundaries' corner points (third

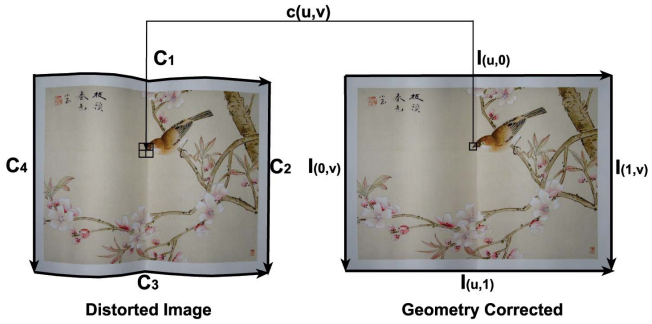

Figure 4: Curves $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4$ correspond to the rectilinear edges of the corrected image $I(u,0), I(1,v), I(u,1), I(0,v)$ respectively. Coons patch interpolation using the $\mathbf{c}_i(\cdot)$ curves provides a mapping between restored (u,v) and the distorted (x,y) coordinates. Geometric correction can performed by pointwise bilinear samples of the (x,y) points. Note that shading correction has not been applied and the restored image may still appear distorted.

term). Because of the need for the corrective function in the interpolant, the function is not considered a bi-linear interpolation, but instead is a bi-linearly "blended" interpolant [7].

The advantage of using equation 5 with four curves, instead of two as defined in the original ruled surface equation 2, is that it is not necessary to specify the true opposite boundary curves. One pair of opposite boundaries should be straight-lines, however, in practice they are not straight due to non-linearites in the image system. The non-linearities are encoded in the boundary splines and their effects will be compensated for in the unwarping process.

2.4. Geometric Distortion Removal

Using the four boundary curves, the function $\mathbf{c}(u,v)$ provides a mapping between (u,v) coordinates in the rectilinear image space $I(u,v)$ to their corresponding (x,y) coordinates in the distorted image. Constructing the restored rectilinear image is performed by pointwise bilinear re-sampling of the distorted image using the relationship $(u,v) \rightarrow (x,y)$. The size of the restored image is specified either by the user or set according to pixel length of the horizontal and vertical c_i curves in the distorted image.

3. Shading Correction

Shading is a strong visual cue for shape. Correcting geometric distortion without addressing shading artifacts can produce restored images that still *appear* perceptually distorted (see Figure 4). For applications, such as text-based imaging with OCR processing, shading artifacts can be ignored, however, for images of materials where the original content is desired (for example, images of artwork), shading artifacts must be removed to produce a perceptually correct image.

3.1. Shading Artifacts Removal

For shading correction, we only consider the luminance component of the image. Our input image is represented in YUV colorspace and the Y-channel (luminance) of the input image, I_Y, can be expressed as the product of the *intrinsic* illumination image, L_Y, and the *intrinsic* reflectance image, R_Y, [2] as follows:

$$I_Y = L_Y \cdot R_Y, \tag{6}$$

where $\cdot$ is a pixel-wise multiple between the two images. If we can successfully compute L_Y, we can derive the intrinsic reflectance image R_Y as:

$$R_Y = e^{logI_Y - logL_Y}. \tag{7}$$

Once we compute the reflectance image R_Y, we can generate an image with uniform shading by multiplying the reflectance image by a constant c, such that $I_Y^{new} = cR_Y$. While the intrinsic image model is view-dependent, we only need to restore our image from a single viewpoint. As a result, this simple intrinsic image model can be used as a reasonable approximation of our imaged scene.

Given a single image, solving for the intrinsic illumination image is ill-posed as the number of unknowns (L and R) is more than the given input I. Several authors have proposed approaches to estimate L from both single and multiple temporal images [8, 14, 16]. For our algorithm, however, we can exploit the fact that almost all printed materials have a uniformly colored (typically white) margin, or border, about the page, void of printed content. The material's reflectance property can be assumed to be the same everywhere on the border. Intensity variations along this uniform border are due to the amount of illumination present on the material's surface. These intensities provide illumination samples of the material at the borders only. The illumination in the interior of the material still needs to be computed. This again becomes a boundary interpolation problem which can be addressed using an approach similar to equation 5.

3.2. Illumination Image Estimation

Let (u,v) represent 2D image coordinates aligned in the restored Y-channel image, I_Y. From Figure 5, we can see that the intensity value at the 2D coordinates (x,y) along the boundary of $\mathbf{c}_1(u)$ should map to $L_1(u)$, i.e. the top edge of the undistorted image. This correspondence holds for other curves, e.g. (x,y) coordinates along the curve $\mathbf{c}_2(v)$ should map to $L_2(v)$. Given the intensity values along the boundaries, the pixel value of interior points can be found using the following equation:

$$\begin{aligned} L_Y(u,v) = &\begin{bmatrix} 1-u & u \end{bmatrix} \begin{bmatrix} L_4(v) \\ L_2(v) \end{bmatrix} \\ &+ \begin{bmatrix} L_1(u) & L_3(u) \end{bmatrix} \begin{bmatrix} 1-v \\ v \end{bmatrix} \\ &- \begin{bmatrix} 1-u & u \end{bmatrix} \begin{bmatrix} L_1(0) & L_2(0) \\ L_3(1) & L_4(1) \end{bmatrix} \begin{bmatrix} 1-v \\ v \end{bmatrix}, \end{aligned} \tag{8}$$

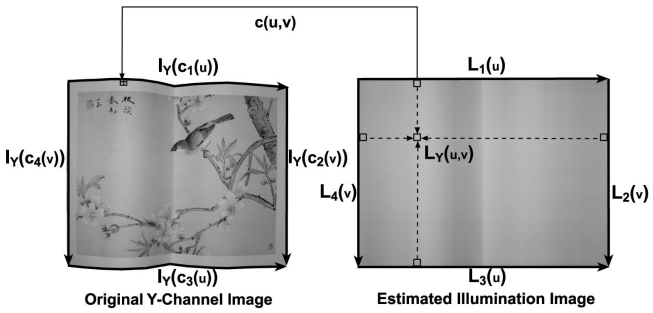

Figure 5: Boundary intensity values along the curves $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4$ correspond to the rectilinear edges of the shadow image $L_1(u), L_2(v), L_3(u), I_4(v)$ respectively. Each inner pixel value of coordinates (x, y) in the shadow image is found from boundary values using the function $L_Y(u, v)$ and pixel is re-sampled using bilinear interpolation.

where

$$L_i(u) = I_Y(\mathbf{c}_i(u)),$$
$$L_i(v) = I_Y(\mathbf{c}_i(v)),$$

where $\mathbf{c}_i(\cdot)$ returns the 2D image coordinate (x, y) along the appropriate boundary curve $\mathbf{c}_i$. The illumination image estimation is also based on Coons patch interpolation, however the data being blended is not (x, y) coordinates but intensity values along the border. Figure 5 shows the results of the illumination estimation. The desired intrinsic reflectance image and the subsequent image with shading artifacts removed can be computed using equation 7.

4. Results

Results of several examples processed by our technique are presented. These examples use the two boundary parameterizations, (1) arc-length and (2) uniform, as mentioned in Section 2.2. Using arc-length parameterization requires no knowledge about the depth change of the imaged material's boundary curves and acceptable results are obtained when the depth distortion is small compared to the size of the imaged material (for example, Figure 1(Top) is corrected with arc-length parameterization). Uniform parameterization is useful to correct item which has significant depth change. For uniform parameterization, a paper checkerboard pattern is placed underneath the imaged material. This pattern undergoes same deformation as the printed media and is used to guide the uniform parameterization. We also show the computed illumination image for our input and show the results of the image with shading artifacts removed. Under controlled imaging environment, automatic border detection can be robustly performed using segmentation. For our experiments we use corner detection of the checkerboard pattern to supply the boundary points $\mathbf{x}_i$ since multiple parameterizations are used for comparison. For illumination estimation, the borders are adjusted slightly to be within the document's margin.

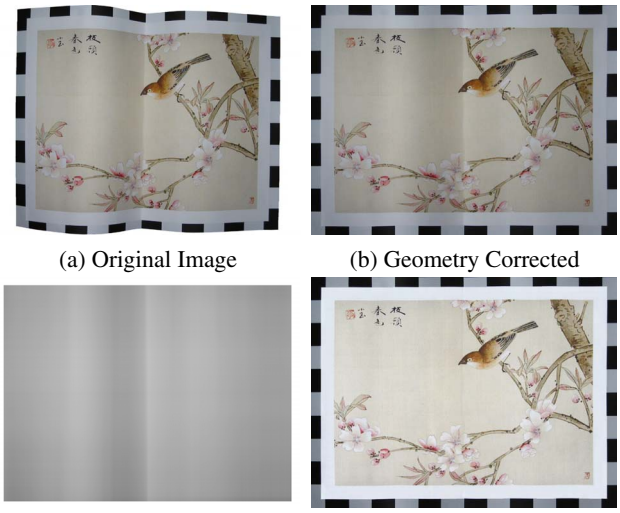

(a) Original Image (b) Geometry Corrected

(c) Estimated Illumination Image (d) Shading Corrected

Figure 6: Example of material with spine distortion. (a) Original image. (b) Corrected image with distortion removed. (c) Estimated intrinsic illumination image. (d) Shading corrected.

4.1. Small Bird Example

Figure 6(a) shows the first example with binder curl distortion. A checkerboard pattern is placed underneath the page and uniform parameterization is used to correct geometric distortion. The result are shown in Figure 6(b). Figure 6(c) shows the estimated illumination image. Figure 6(d) shows the restored image with uniform illumination.

4.2. Buddhist Fold-Out Example

Figure 7(a) shows an example which is a large fold-out page from an oversized art book. This is an excellent representation of materials where existing techniques are not applicable. This book has many pages that fold-out to display very wide content. The page is $80cm \times 33cm$ in width and height and is very difficult to image completely flat and exhibits both fold and curl distortion.

Figure 7-I(b) shows the result using arc-length parameterization. The corners of the checkerboard pattern are used as the 2D feature points to define the boundary splines. Notice that the corrected image has artifacts due to the depth distortion. This can be verified by looking at the uneven black and white square pattern along the boundaries. Figure 7-II(b) shows the result which parameterized the boundary using the same 2D feature points but with uniform parameterization. This parameterization is justified because the 2D features are from uniform samples on the 3D curve obtained from the inserted pattern. The depth distortion has been corrected. Figure 7(c) shows the estimated illumination images using white border information from the input boundary. Figure 7(d) shows the new images under uniform illumination.

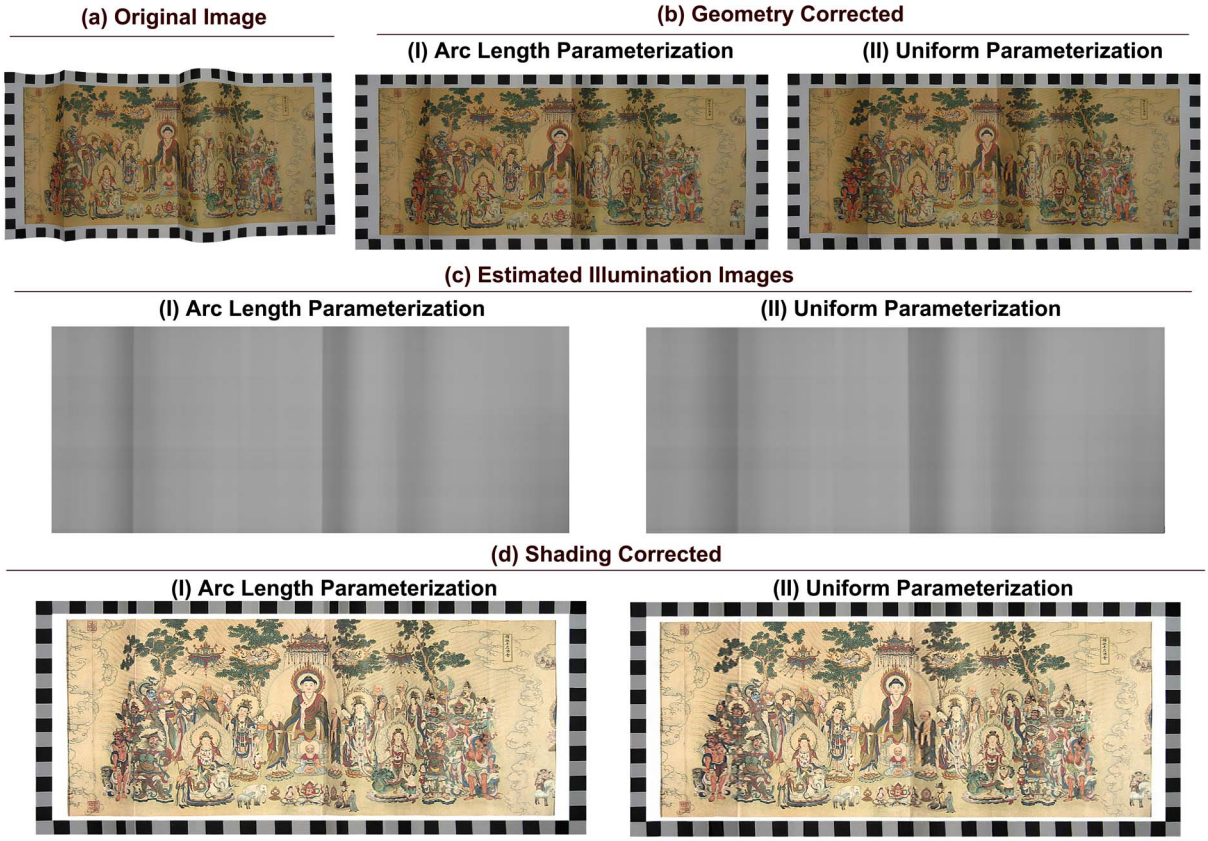

Figure 7: Example of wide fold-out page from an art book, exhibiting both binder-curl and fold distortion. Corrected images using arc-length parameterization are shown on column I, those using uniform parameterization is shown in column II. The images are labeled: (a) Original image. (b) Geometry corrected. (c) Estimated illumination images. (d) Shading corrected.

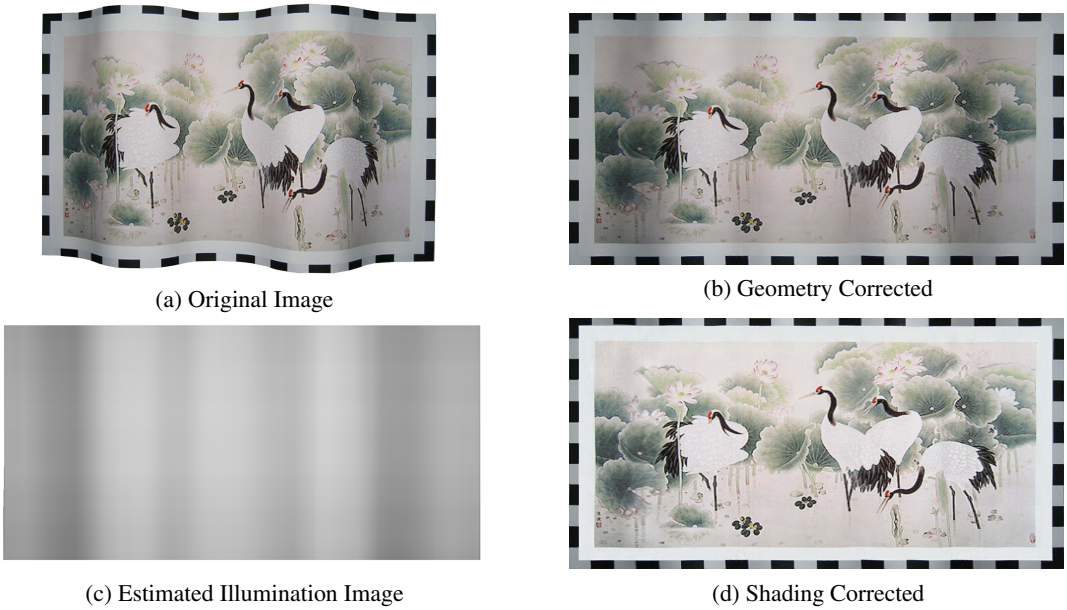

Figure 8: Example of material with roll (curl) distortion. 2.5D information is incorporated to guide the uniform parameterization. (a) Original image. (b) Geometry corrected using uniform boundary sampling. (c) Estimated illumination image. (d) Shading corrected.

4.3. Birds Example

Figure 8(a) shows another example of imaged art with significant "roll" distortion. Uniform parameterization is used to correct the geometric distortion, shown in 8(b). The contents of the imaged material have been rectified. Figure 8(c-d) shows the derived illumination image and resulting image with shading corrected.

5. Conclusion

This paper presents two innovations for document imaging. First, we show that the printed materials can be modeled using a ruled surface composed of opposite boundary curves. This general surface model can model several common distortion found in printed materials, including: skew, folds, and binder-curl. We show that the projection of this model can be parameterized using Coons patch blending with four boundary curves. This 2D interpolation provides a mapping between the distorted and restored image allowing shape distortion to be corrected. In addition, 2.5D information can be easily incorporated into the interpolation process by changing only the curve's parameterization. This allows the correction of materials with significant depth distortion.

Secondly, we showed that the same interpolation framework can be used to estimate the intrinsic illumination image of the input image. From this illumination image, we can generate a new image with shading artifacts removed.

Currently, we use a physical pattern to guide the uniform parameterization. The use of such a pattern is reasonable in a restoration context. Our personal experience with digitization efforts, for preservation and archival purposes, is that a great deal of care and time is spent on each imaged item and additional props (such as rulers and color strips) are routinely incorporated into the imaging environment. Using a paper pattern under the imaged material would not present a great burden and could be reasonably adopted, especially if it is to facilitate restoration. We are currently working on techniques to remove the need for this pattern. Using the illumination image obtained from arc-length parameterization, we can apply shape from shading techniques to get an estimation of the 3D arc-length near the boundary to parameterize the 2D curve without the need of a pattern.

In summary, we have presented a novel approach based on boundary interpolation that can correct geometric distortion and shading artifacts present in images of printed materials. Our algorithm can simultaneously correct a variety of geometric distortions, including skew, folding, binder curl. In addition, the same interpolation framework can be used to estimate the illumination component of the input image which is used to correct shading artifacts.

6. Acknowledgments

We gratefully acknowledge grants DAG02/03.EG02 and HKUST6177/03E for supporting this work.

References

[1] Avanidra and S. Chaudhuri. Robust detection of skew in document images. *IEEE Transactions on Image Processing*, 6(2):344–349, 1997.

[2] H.G. Barrow and J.M. Tenenbaum. *Recovering intrinic scene characterisitcs from images*. Academic Press, 1978.

[3] M. S. Brown and W. B. Seales. Document restoration using 3D shape. In *ICCV '01*, July 9-12 2001.

[4] H. Cao, X. Ding, and C. Liu. A cylindrical surface model to rectify the bound document image. In *ICCV'2003*, pages 228–233, 2003.

[5] H. Cao, X. Ding, and C. Liu. Rectifying the bound document image captured by the camera: A model based approach. In *Proc. 7th International Conference on Document Analsysis and Recognition*, 2003.

[6] S. Coons. Surfaces for computer aided design. Technical report, MIT, 1968. Available as AD 663 504 from the National Technical Information service, Springfield, VA 22161.

[7] G. Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, San Diego, CA, 1990.

[8] B.V. Funt, M.S. Drew, and M. Brockington. Recovering shading from color images. In *ECCV'92*, pages 123–132, 1992.

[9] B. Gatos, N. Papamarkos, and C. Chamzas. Skew detection in text line position determination in digitized documents. *Pattern Recognition*, 30(9):1505–1519, 1997.

[10] H. Jiang, C. Han, and K. Fan. Automated page orientation and skew angle detection for binary document images. *Pattern Recognition Letters*, 18(7):125–133, 1997.

[11] T. Kanungo, R. Haralick, and I. Phillips. Global and local document degradation models. In *ICDAR-93*, pages 730–734, Tsukuba, Japan, October 20-22 1993.

[12] U. Pal and B. Chaudhuri. An improved document skew angle estimation technique. *Pattern Recognition Letters*, 8(17):899–904, July 1996.

[13] M. Pilu. Undoing paper curl distortion using applicable surfaces. In *CVPR '01*, Dec 11-13 2001.

[14] M. Tappen, W.T. Freeman, and E. Adelson. Recovering intrinsic images from a single image. In *Advances in Neural Information Processing Systems 15 (NIPS)*. MIT Press, 2003.

[15] T. Wada, H. Ukida, and T. Matsuyama. Shape from shading with interreflections under proximal light source. In *ICCV '95*, pages 66–71, 1995.

[16] Y. Weiss. Deriving intrinsic images from image sequences. In *ICCV'01*, pages II: 68–75, July 2001.

[17] Z. Zhang and C. L. Tan. Restortation of images scanned from thick bound documents. In *ICIP '01*, Thessaloniki, Greece, Oct 2001.