Photometric and Geometric Restoration of Document Images Using Inpainting and Shape-from-Shading

Li Zhang
School of Computing
3 Science Drive 2
National University of Singapore

Andy M. Yip
Department of Mathematics
2 Science Drive 2
National University of Singapore

Chew Lim Tan
School of Computing
3 Science Drive 2
National University of Singapore

Abstract

The popularity of current hand-held digital imaging devices such as camera phones, PDAs, camcorders has promoted the use of digital cameras to capture document images for daily information recording purpose. However, the captured images often contain photometric and geometric distortions when the documents are of non-planar shapes, which cause significant problems to various document image analysis (DIA) tasks such as OCR. In this paper, we propose a restoration framework that removes both photometric and geometric distortions in smoothly warped document images to facilitate human perception and machine recognition. First, the photometric distortions are corrected by separating the shading image from the reflectance image using inpainting and surface fitting techniques. Next, a 2-pass Shape-from-Shading (SFS) method is exploited to recover the document's surface shape based on the extracted shading image. Once the document's shape is obtained, the geometric distortions are rectified through a physically-based flattening process. Experiments on real document images show the performance of each sub-task and demonstrate a complete solution to the restoration of physically-distorted document images.

Introduction

The increasing popularity of current digital cameras has brought about a new trend of digitizing physical documents especially for those deteriorated ancient manuscripts. The contactless imaging property preserves the fidelity of the documents but inevitably introduces certain distortions due to the non-planar geometric shape of the document and the perspective projection principle. In order to successfully apply camera imaging to the digitization of fragile non-planar documents, normalization methods need to be developed to correct all these distortions for subsequent machine recognition tasks. Apart from digitizing historical documents, camera imaging also provides a simple way to record daily information. This is especially true with the fast emerging hand-held digital imaging devices in current cell phones, PDAs, camcorders, etc. The high portability, versatile functionality and low pricing properties of such devices have promoted their use as personal photocopiers which can be carried everywhere with ease (Doermann, Liang, & Li 2003). The successful application of camera imaging in these areas requires a comprehensive study on camera-based document image processing techniques including text detection and extraction, normalization, enhancement and recognition. In particular, the restoration of photometric and geometric distortions is a necessary pre-processing step that ensures the smooth execution of the subsequent DIA tasks.

Related Work

Methods to correct geometric distortions such as rolling curls, binding curls and folds have been widely studied in the past few decades for both scanned or camera-based document images. Generally speaking, they can be categorized as either 2D-based or 3D-based, depending on whether the document's 3D geometry is used in the restoration process. Most 2D-based approaches use image processing techniques to find geometric transformations between the input and the output image. Each point in the output image is mapped to a corresponding point in the input image through a spatial transformation function, such as affine, perspective, bilinear or polynomial. To find such geometric transformations, the 2D textual content is often needed, such as reference points (Tang & Suen 1993), text lines (Zhang & Tan 2003), document boundaries (Brown & Tsoi 2006), etc. On the other hand, in order to get a more accurate representation of the physical warping and also to handle more complex warping distortions, efforts have been put into the acquisition of the document's 3D geometry. In particular, some researchers try to capture the surface shape using special setups such as structured lighting (Pilu 2001; Brown & Seales 2004; Brown & Pisula 2005), laser scans (Chua *et al.* 2005), or CT scans for opaque objects (Lin & Seales 2005). The 3D model is then mapped to a plane using various numerical techniques. Moreover, methods are also proposed to reconstruct the document's shape based on one or more 2D images. Most of these approaches rely on assumptions of some special properties of the surface shape (Cao, Ding, & Liu 2003) or the imaging devices such as flatbed scanners (Zhang *et al.* 2005), stereo vision systems (Yamashita *et al.* 2004), or video cameras (Iketani, Sato, & Ikeda 2006).

Photometric restoration in terms of shading correction can be easily done using simple binarization methods on

scanned document images (Zhang & Tan 2003). Nevertheless, methods are also developed to correct scanned shadings based on reverse SFS formulation (Zhang *et al.* 2005). Due to the intricate lighting environment, shading artifacts in camera-based images are more complicated to handle. Brown and Tsoi propose to correct shading distortions in art materials using boundary interpolation borrowed from the corresponding geometric correction approach (Brown & Tsoi 2006). Sun *et al.* introduce a method that corrects the shading by classifying the intensity changes into illumination changes or reflectance changes (Sun *et al.* 2005) based on the notion of intrinsic images (Barrow & Tenenbaum. 1978). Outside the document image domain, more studies have been focusing on the extraction of intrinsic images on real-scene images (Olmos & Kingdom 2004; Tappen, Freeman, & Adelson 2005). Despite all the efforts put into deriving the intrinsic images, there is no single exact solution because the decomposition of the intensity image into its two intrinsic components is theoretically not unique.

Our Contribution

In this paper, we present a restoration framework that corrects both photometric and geometric distortions in smoothly warped document images. To correct photometric distortions, we first make use of digital inpainting and surface fitting techniques to extract a smooth shading image which is then separated from the original intensity image to obtain the reflectance image. Once the shading image is extracted, it can also be used to reconstruct the 3D shape based on the SFS concept. Here we propose a 2-pass SFS method that consists of a viscosity framework and a minimization scheme with regularization techniques to ensure a good approximation of the document's surface shape. Next, we can map the original 2D image to the reconstructed shape and remove the geometric distortions by physically flattening the 3D model using numerical simulation. Alternatively, we can use the photometrically restored image as texture and obtain a final image that has both photometric and geometric distortions removed. A general work flow of the restoration framework is shown in Figure 1.

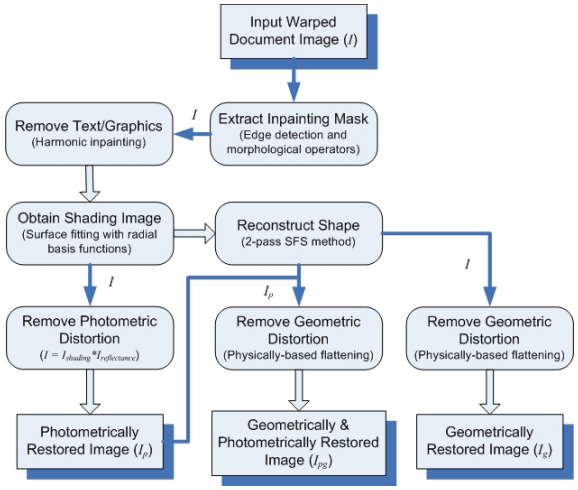

Figure 1: A general work flow of the document image restoration framework.

Shading Image Extraction

Shading image extraction is a very important step that affects both the photometric and the geometric correction processes significantly. Essentially, the shading image should contain all the shading artifacts caused by changes of the surface normal with respect to the illumination direction. The quality of the shading image thus determines how well the shading artifacts are detected and removed. On the other hand, the performance of the SFS algorithm also greatly relies on the accuracy of the extracted shading image. Assuming the given document image has a uniform-colored background such as the normal printed notes, book pages, etc, an effective cue for differentiating shading from reflectance is the printed regions. It has been observed that luminance variations accompanied by color changes are usually variations in reflectance while luminance variations unaccompanied by color changes are variations in illumination (Olmos & Kingdom 2004). Therefore, the first step is to identify the text/graphics locations.

Inpainting Mask Generation

We first make use of an edge-based method to identify pixels that are of high contrast to the background. Next, morphological operations are applied to the edge-detected image to generate a mask for the printed regions. The detailed procedures are as follows: 1) Convert color images into gray-scale by picking the luminance component of a color model such as the V-component of the HSV model or the I-component of the HSI model; 2) Detect edges using canny edge detector with post-processings such as non-maximum suppression and streaking elimination; 3) Perform morphological dilation followed by closing. An example of the inpainting mask generated is shown in Figure 2 (b).

Harmonic/TV Inpainting

Once the mask of the printed regions is generated, an inpainting technique is used to fill in the masked regions based on the neighboring background pixels. This is essentially to recover the shading in the printed regions based on the assumption that the local variation of shading is small. Digital inpainting was pioneered by Bertalmio *et al.* (Bertalmio *et al.* 2000) and has since been applied to a variety of image processing applications. Here we use digital inpainting as a way of recovering the shading. In particular, we experimented with two non-texture inpainting models, harmonic and TV inpainting (Chan & Shen 2002). The main objective is to find I that minimizes the following energy in a continuous domain Ω:

$$E(I) = \int_{\Omega} \chi \cdot (I - I_0)^2 dx + \lambda \int_{\Omega} |\nabla I|^2 dx \qquad (1)$$

where $\lambda > 0$ is a smoothness parameter and χ denotes the characteristic function:

$$\chi(x) = \begin{cases} 1, & x \in \Omega \backslash \Omega_H \\ 0, & \text{otherwise} \end{cases} \qquad (2)$$

For TV inpainting, instead of using a penalty term $\int |\nabla I|^2 dx$ in Eq. 1, which is infinite for discontinuous functions, we use $\int |\nabla I| dx$ which allows discontinuous functions as minimizers. An example of the inpainted image is shown in Figure 2 (c). Both harmonic and TV inpainting are essentially local models, in which the inpainting is determined by the existing information in the vicinity of the inpainted domain. The main difference is that harmonic inpainting builds very smooth solutions and thus does not cope well with edges, while TV inpainting is able to restore narrow broken smooth edges which often exist in document images due to overlaid texts.

Surface Fitting with RBF

The initial estimation of the shading after inpainting may contain pepper noises due to errors in the extracted mask. One way to solve this problem is to iteratively improve the mask until the extracted shading is satisfactory. Alternatively, we can remove the pepper noises by using a surface fitting algorithm with radial basis functions (RBF) (Carr *et al.* 2003). This is especially useful when the shading needs to be smooth for further surface reconstruction tasks. Typically, given a set of 3D points $\{(\mathbf{x}_i, f(\mathbf{x}_i)),\ i = 1, 2, \cdots, m\}$, a fitted surface can be expressed as:

$$g(x) = \sum_{j=1}^{n} \alpha_j h(\mathbf{x} - \mathbf{y}_j) \tag{3}$$

where $\{\mathbf{y}_j,\ j = 1, 2, \cdots, n\}$ is a set of collocation points and $h(\mathbf{x})$ is the radial basis function. The goal is to find the coefficients α_j that minimize the least squares error:

$$e = \min_{\alpha_1, \cdots, \alpha_n} \left\{ \sum_{i=1}^{m} (g(\mathbf{x}_i) - f(\mathbf{x}_i))^2 \right\} \tag{4}$$

with optional boundary conditions.Various kernel functions of different smoothness can be used. Here we use Multiquadrics: $h(\mathbf{x}) = \sqrt{||\mathbf{x}||^2 + c^2}$, where c is a constant. The advantages of using RBF fitting are: 1) It gives explicit formula for derivatives which are more accurate and less noisy than finite difference; 2) It is also easy to incorporate various types of boundary conditions; 3) Unlike polynomial fitting, RBF is more flexible and can be used to fit more complicated surfaces. Figure 2 (d) shows an example of the smooth shading image extracted using surface fitting.

Photometric Restoration

Once the shading image is extracted, it is easy to derive the reflectance image based on the notion of intrinsic images. For Lambertian surfaces, the intensity image is the product of the shading image and the reflectance image (Barrow & Tenenbaum. 1978). Consider the luminance component of the HSV model, we have $I = I_s \cdot I_r$. Now given the shading I_s, the reflectance image I_r can be computed as: $I_r = e^{log\, I - log\, I_s}$. The photometrically restored image can thus be computed as: $I_p = k \cdot I_r$, where $k \in [0, 1]$.

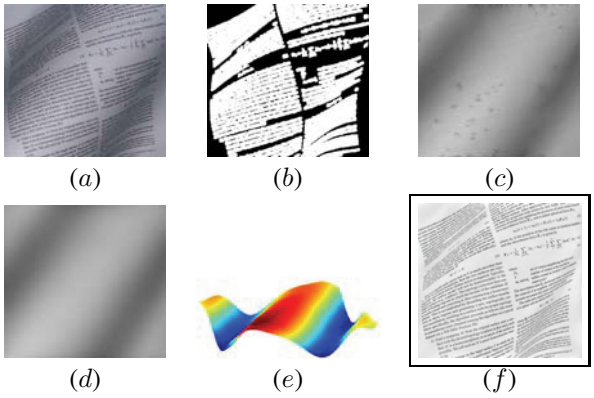

Figure 2: (a) Image of an arbitrarily warped document page; (b) Inpainting mask; (c) One-pass inpainted image; (d) Shading image using RBF fitting; (e) Fitted 3D surface of the shading image; (f) Photometrically restored image.

Geometric Restoration

To get an accurate representation of the geometric warping for better restoration results, we try to reconstruct the document's surface shape. Given the shading image, this can be done using SFS techniques that aim to reconstruct shapes using shading variations in a single image. Research in this field was pioneered by Horn who first formulated the SFS problem as to find the solution of a nonlinear first-order PDE called the brightness equation (Winston 1975). Following this, a series of variational methods (Ikeuchi & Horn 1981; Horn & Brooks 1986; Frankot & Chellappa 1988) are developed which try to minimize an energy function that often comprises of an integral of the brightness error to find the solution. Later Oliensis and Dupuis propose to cast SFS as an optimal control problem and directly find the surface depth map (Oliensis & Dupuis 1991). This leads to a new set of propagation approaches based on the theory of viscosity solutions to Hamilton-Jacobi equations (Osher & Sethian 1988; Rouy & Tourin 1992; Prados, Faugeras, & Rouy 2002). More comprehensive surveys can be found in (Zhang *et al.* 1999; Durou, Falcone, & Sagona 2004).

Perspective SFS formulation

Here we focus on the SFS problem under the assumption of perspective projection with a close point light source simulating the camera's flash. Since the distance between the camera and the object is usually much greater than the focal length, it is reasonable to assume that the light source is aligned with the optical center. Suppose the document surface follows Lambertian reflection, we have the image irradiance equation as follows:

$$I(u, v) = \frac{-pu - qv + f}{\sqrt{p^2 + q^2 + 1}\sqrt{u^2 + v^2 + f^2}} \tag{5}$$

where (u, v) is the normalized image coordinates with respect to the principle component coordinate (u_0, v_0), $I(u, v)$ is the corresponding pixel intensity, $(-p, -q, 1)$ is the surface normal and f is the focal length. Let $z(u, v)$ denote the

distance from the surface point (x, y) to the u-v plane, we then have the surface gradients: $p = \frac{\partial z}{\partial u}$ and $q = \frac{\partial z}{\partial v}$.

Pass I: Viscosity Framework

We observe that the image irradiance equation in Eq. 5 can be written in the form of a static HJ equation:

$$\begin{cases} H(u, v, \nabla z) = R(u, v), & (u, v) \in \Omega \\ z(u, v) = B(u, v), & (u, v) \in \Gamma \subset \Omega \end{cases} \tag{6}$$

where Ω denotes the image plane, Γ denotes a set of points in Ω where the value of $z(u, v)$ is known to be $B(u, v)$, called the boundary values, though they may be located in the interior of Ω, and also:

$$\begin{cases} H(u, v, \nabla z) & = I\sqrt{p^2 + q^2 + 1}\sqrt{u^2 + v^2 + f^2} \\ & \quad + pu + qv - f \\ R(u, v) = 0 \end{cases} \tag{7}$$

With the formulation in Eq. 7, we can then exploit an iterative sweeping strategy (Tsai *et al.* 2003) to solve for $z(u, v)$ with an update formula based on the Lax-Friedrichs Hamiltonian (Kao, Osher, & Qian 2004) given as:

$$\begin{aligned} z_{u,v}^{n+1} &= \frac{1}{\frac{\sigma_u}{\Delta u} + \frac{\sigma_v}{\Delta v}}(R(u, v) - H(p, q) + \sigma_u u_m + \sigma_v v_m) \\ p &= \frac{z_{u+1,v} - z_{u-1,v}}{2\Delta u} \quad q = \frac{z_{u,v+1} - z_{u,v-1}}{2\Delta v} \\ u_m &= \frac{z_{u+1,v} + z_{u-1,v}}{2\Delta u} \quad v_m = \frac{z_{u,v+1} + z_{u,v-1}}{2\Delta v} \end{aligned} \tag{8}$$

where $(\Delta u, \Delta v)$ is the grid size, σ_u and σ_v are artificial viscosities satisfying $\sigma_u \geq \max_{u,v,p,q} |\frac{\partial H}{\partial p}|$ and $\sigma_v \geq \max_{u,v,p,q} |\frac{\partial H}{\partial q}|$. In particular, for Eq. 7, we let

$$\begin{aligned} \sigma_u &= \max_{u,v,p,q} \left|\frac{\partial H}{\partial p}\right| = \max_{u,v}\{\max\{|I_p + u|, |I_p - u|\}\} \\ \sigma_v &= \max_{u,v,p,q} \left|\frac{\partial H}{\partial q}\right| = \max_{u,v}\{\max\{|I_p + v|, |I_p - v|\}\} \end{aligned} \tag{9}$$

where $I_p = I\sqrt{u^2 + v^2 + f^2}$.

The iterative sweeping strategy is based on a fast sweeping scheme (Tsai *et al.* 2003). First, the surface is initialized with the boundary values $B(u, v)$. Next, the value of $z(u, v)$ is updated by sweeping through the grid in four alternating directions. Finally, after each sweep, the height values are evaluated at the four boundaries where the update formula fails to compute. The complexity of fast sweeping is $O(N)$ where N is the number of grid points.

Pass II: Minimization with regularization

Essentially, Pass I gives a viscosity solution to the SFS problem. However, when applying to real camera images, the noises in the original image or the approximated shading image often cause abrupt errors in the reconstructed result. In order to improve the OCR performance on those text-dominant images, a better reconstruction is necessary. Therefore, we further apply a least squares method with a regularization term to smooth out the abrupt ridges caused by noises or errors in the approximated shading. On the other hand, the result in Pass I also provides a good initialization for the minimization method, which avoids the problem of being trapped in local minima. The minimization method is based on the variational SFS formulation discussed in (Horn & Brooks 1986; Crouzil, Descombes, & Durou 2003), with the energy:

$$\begin{aligned} F_1(p, q) = &\iint_\Omega [I(p, q) - E(u, v)]^2 \, dudv \\ &+ \lambda_i \iint_\Omega \left[\frac{\partial p}{\partial v} - \frac{\partial q}{\partial u}\right]^2 dudv \\ &+ \lambda_s \iint_\Omega \left[|\nabla p|^2 + |\nabla q|^2\right] dudv \end{aligned} \tag{10}$$

where $I(p, q)$ is the image irradiance defined in Eq. 5, $E(u, v)$ is the image intensity, λ_i and λ_s are the integrability and smoothing coefficients, respectively. Similarly, to derive the height z from p and q, we use the second energy:

$$F_2(z) = \iint_\Omega \left[\left(\frac{\partial z}{\partial u} - p\right)^2 + \left(\frac{\partial z}{\partial v} - q\right)^2\right] dudv \tag{11}$$

To numerically minimize the above two energy functions, we minimize their discrete counterparts using the steepest descent method with a simple line search with Armijo condition (Fletcher 1987). The partial derivatives of p and q are discretized using forward finite difference. In order to consider all the boundary points, we need to enforce a Neumann boundary condition on p and q. Suppose (m, n) is the size of D_Ω, we then define $p_{0,\cdot} = p_{1,\cdot}$, $p_{m,\cdot} = p_{m+1,\cdot}$, $p_{\cdot,0} = p_{\cdot,1}$ and $p_{\cdot,n} = p_{\cdot,n+1}$. Same applies to q. Given an initialization of (p, q) from Pass I, we first apply an iterative process to find a better configuration $(p, q)_m$ that minimizes F_1 within a certain number of iterations. Next, this new configuration will be used in the evaluation of F_2 in which z is initialized as the result of Pass I. Similarly, we obtain a new configuration z_m which is the final result.

Physically-based Flattening

Once the surface shape of the document is obtained, the restoration can be done by flattening the warped shape to a plane through a numerical method (Brown & Seales 2004; Chua *et al.* 2005). Before this, a 2D texture image needs to be mapped onto the reconstructed shape so that it will be restored accordingly when the shape is flattened. This is done by mapping the 3D coordinates of the mesh points to their 2D image coordinates based on the perspective projection principle given as: $x/u = y/v = z/f$. In addition, if the photometrically restored image is used as the texture, we obtain the final image with both geometric and photometric distortions removed.

Experimental Results

All the warped images in our experiments are taken in a relatively dark environment with the camera's flash acting as a close point light source. The camera's focal length and

principle components are obtained through a simple calibration procedure. Typically, for an image of size 1600×1200, we have $f = 1348.28$ and $(u_0, v_0) = (790.24, 581.85)$ in pixel size. The images are cropped to avoid lens distortions near the corners. Gamma correction is performed on the extracted shading images when necessary.

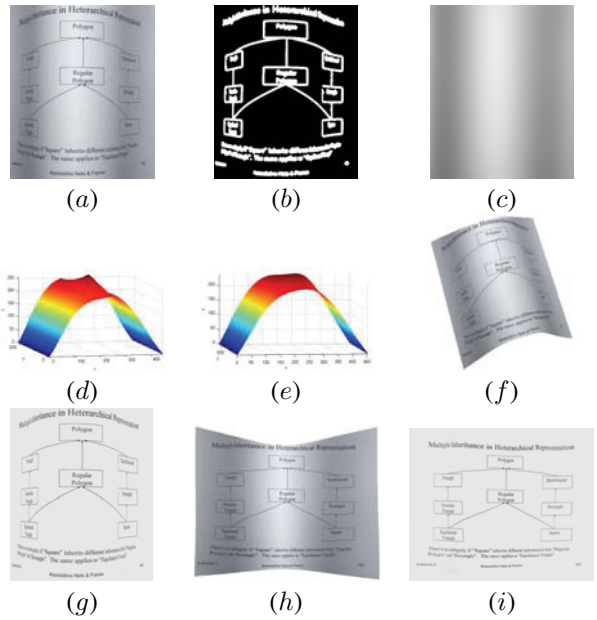

Figure 3: (a) Original warped image (cropped from a 1600×1200 image); (b) Inpainting mask; (c) Extracted shading image after surface fitting; (d) Reconstructed shape after pass I; (e) Shape after pass II; (f) Surface mesh with texture; (g) Photometrically restored image with $k = 0.9$; (h) Geometrically restored image; (i) Final restored image.

Figure 3 shows the results on an arbitrarily curved document with mixed figures and texts. From Figure 3 (c) and (g), we can see that the extracted shading image reflects the illumination change nicely and is separated well from the reflectance image. Next, Figure 3 (d) and (e) demonstrate how the minimization procedure improves the reconstructed shape through the second pass. The first pass is initialized by setting the height at the two vertical boundaries to 0. This could be any arbitrary value because the reconstructed shape is invariant up to a translation factor. Figure 3 (f) shows the reconstructed shape with the original warped image mapped as the texture. Finally, Figure 3 (h) and (i) give the geometrically restored image and the final image with shadings removed, respectively. It is noticed that the restored image is much better improved comparing to the original image although there are still some distortions due to the imperfection of the estimated shape. Better results are expected if a more accurate light source location is provided.

Figure 4 shows another example of a diagonally curved document with mainly text contents. This randomly curved document does not have obvious boundaries lying on the same plane. In this case, we use the singular points as the initialization condition in the first pass. Experiments show

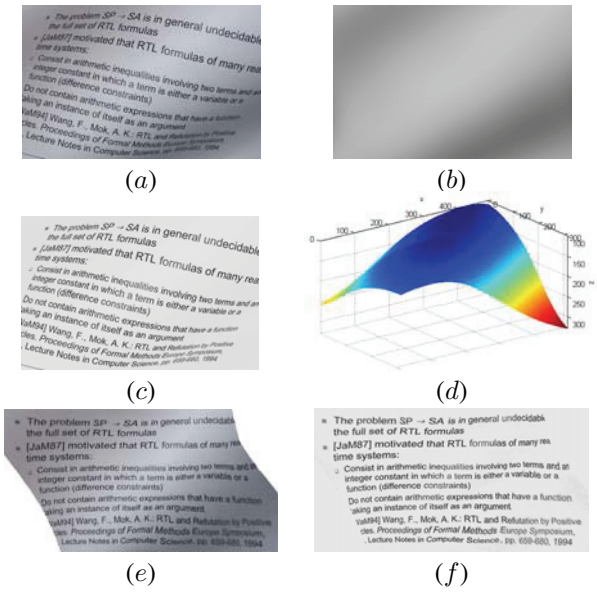

Figure 4: (a) Original warped image; (b) Extracted shading image; (c) Photometrically restored image with $k = 0.9$; (d) Reconstructed shape by initializing singular points; (e) Geometrically restored image; (f) Final restored image.

that even if the singular points are slightly off, the result is not affected much. Figure 4 (d) and (e) show the reconstructed shape and the geometrically restored image, respectively. We can see that the shape does emulate the original curvature though it is not a perfect reconstruction. The restored image also shows a better visual appearance despite some unremoved distortions. In terms of OCR performance, the restored image gives a word precision of 95.6% comparing to 38.9% on the original image. Moreover, we have collected a total of 20 warped document images with approximately 2,400 words for OCR testing. The average word precision is 94.3% on the restored images in contrast to 43.6% on the original images.

Conclusion

In this paper, we present a restoration framework for correcting both photometric and geometric distortions in smoothly warped document images captured using hand-held digital cameras. The photometric restoration is based on the notion of intrinsic images, which first extracts a smooth shading image through inpainting and surface fitting techniques and then separates it from the reflectance image to remove the shading distortions. The geometric restoration also makes use of the extracted shading to recover the document's surface shape in order to obtain an accurate representation of the physical warping. By flattening the reconstructed 3D geometry with the corresponding 2D texture image, the final restored image is obtained with both photometric and geometric distortions removed. The whole framework is tested on various real document images and shows encouraging results. Although the restored images are not perfect, there are still rooms for improvement such as by relaxing the assump-

tion on the light source location and enhancing the shading estimation function, etc. Moreover, this framework gives a good start to restorations based on a single 2D image.

Acknowledgment

This project is supported by URC grant R252-000-202-112 and A*STAR grant 0421010085.

References

Barrow, H., and Tenenbaum., J. 1978. *Recovering Intrinsic Scene Characteristics from Images*. Academic Press.

Bertalmio, M.; Sapiro, G.; Ballester, C.; and Caselles, V. 2000. Image inpainting. *SIGGRAPH 2000* 1:417C424.

Brown, M., and Pisula, C. 2005. Conformal deskewing of non-planar documents. *CVPR'05* 1:998–1004.

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

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

Cao, H.; Ding, X.; and Liu, C. 2003. A cylindrical model to rectify the bound document image. *IEEE Int'l Conf. on Computer Vision* 2:228–233.

Carr, J.; Beatson, R.; McCallum, B.; Fright, W.; McLennan, T.; and Mitchell, T. 2003. Smooth surface reconstruction from noisy range data. *Graphite 2003* 119 –297.

Chan, T., and Shen, J. 2002. Mathematical models for local nontexture inpaintings. *SIAM Journal on Applied Mathematics* 62(3):1019–1043.

Chua, K.; Zhang, L.; Zhang, Y.; and Tan, C. 2005. A fast and stable approach for restoration of warped document images. *8th Int'l Conf. on Document Analysis and Recognition* 1:384–388.

Crouzil, A.; Descombes, X.; and Durou, J. 2003. A multiresolution approach for shape from shading coupling deterministic and stochastic optimization. *IEEE Trans. on PAMI* 25(11):1416–1421.

Doermann, D.; Liang, J.; and Li, H. 2003. Progress in camera-based document image analysis. *7th Int'l Conf. on Document Analysis and Recognition* 606–616.

Durou, J.; Falcone, M.; and Sagona, M. 2004. A survey of numerical methods for shape from shading. Research Report from IRIT.

Fletcher, R. 1987. *Practical Methods of Optimization*. New York: John Wiley and Sons, 2nd edition.

Frankot, R., and Chellappa, R. 1988. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. on PAMI* 10:439–451.

Horn, B., and Brooks, M. 1986. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing* 33(2):174–208.

Iketani, A.; Sato, T.; and Ikeda, S. 2006. Video mosaicing for curved documents based on structure from motion. *18th Int'l Conf. on Pattern Recognition* 1:391–396.

Ikeuchi, K., and Horn, B. 1981. Numerical shape from shading and occluding boundaries. *Journal of Artificial Intelligence* 17(1):141–184.

Kao, C.; Osher, S.; and Qian, J. 2004. Lax-friedrichs sweeping scheme for static hamilton-jacobi equations. *Journal of Computational Physics* 196(1):367–391.

Lin, Y., and Seales, W. 2005. Opaque document imaging: Building images of inaccessible texts. *IEEE Int'l Conf. on Computer Vision* 1:662–669.

Oliensis, J., and Dupuis, P. 1991. Direct method for reconstructing shape from shading. *SPIE* 1:116–128.

Olmos, A., and Kingdom, F. 2004. A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception* 33(12):1463–1473.

Osher, S., and Sethian, J. 1988. Fronts propagating with curvature dependent speed: Algorithms based on the hamilton-jacobi formulation. *Journal of Computational Physics* 79:12–49.

Pilu, M. 2001. Undoing page curl distortion using applicable surfaces. *Int'l Conf. on Computer Vision and Pattern Recognition* 1:67–72.

Prados, E.; Faugeras, O.; and Rouy, E. 2002. Shape from shading and viscosity solutions. *European Conf. on Computer Vision* 1:709–804.

Rouy, E., and Tourin, A. 1992. A viscosity soluition approach to shape from shading. *SIAM Journal on Numerical Analysis* 3(29):867–884.

Sun, M.; Yang, R.; Yun, L.; Landon, G.; Seales, B.; and Brown, M. 2005. Geometric and photometric restoration of distorted documents. *ICCV'05* 2:17–21.

Tang, Y., and Suen, C. 1993. Image transformation approach to nonlinear shape restoration. *IEEE Trans. on Systems, Man, and Cybernetics* 23(1):155–171.

Tappen, M.; Freeman, W.; and Adelson, E. 2005. Recovering intrinsic images from a single image. *IEEE Pattern Recognition and Machine Intelligence* 27(9):1459–1472.

Tsai, Y.; Cheng, L.; Osher, S.; and Zhao, H. 2003. Fast sweeping algorithms for a class of hamilton-jacobi equations. *SIAM J. on Numerical Analysis* 41(2):659–672.

Winston, P. 1975. *The Psychology of Computer Vision : Obtaining shape from shading information (B.K.P. Horn)*. New York: McGraw-Hill.

Yamashita, A.; Kawarago, A.; Kaneko, T.; and Miura, K. 2004. Shape reconstruction and image restoration for nonflat surface of document with a stereo vision system. *17th Int'l Conf. on Pattern Recognition* 1:482–485.

Zhang, Z., and Tan, C. 2003. Correcting document image warping based on regression of curved text lines. *7th Int'l Conf. on Document Analysis and Recognition* 1:589–593.

Zhang, R.; Tsai, P.; Cryer, J.; and Shah, M. 1999. Shape from shading: A survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 21(8):690–706.

Zhang, L.; Zhang, Z.; Tan, C.; and Xia, T. 2005. 3d geometric and optical modeling of warped document images. *IEEE Conf. on CVPR'05* 1:337–342.