# Conformal Deskewing of Non-Planar Documents

Michael S. Brown
C.S.U.M.B.
Monterey Bay, CA

Charles J. Pisula
Apple
Cupertino, CA

## Abstract

*This paper presents an approach that uses conformal mapping to parameterize a document's 3D shape to a 2D plane. Using this conformal parametrization, a restorative mapping between an image of the distorted document and a "flattened" representation of the document can be computed and used to deskew the image. Our experiments show that arbitrarily distorted documents can be restored to within a single pixel of their true planar format. In addition, surface points can be constrained to mapped to specified locations in the restored 2D plane.*

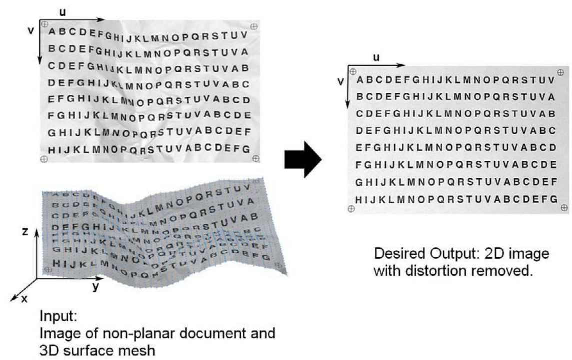


Figure 1: Problem Addressed: The input is (1) an image of a non-planar document and (2) a 3D reconstruction the document's surface. The goal is to produce a *restored* version of the document that appears flat.

## 1 Introduction

Institutions actively digitize printed materials. While such digitization has traditionally been performed using flatbed scanners, an increasingly common trend is to use cameras [8]. Camera-based approaches, however, do little to guarantee that items are flat when imaged. As a result, images of non-planar documents can appear distorted. Techniques to remove distortion from imaged documents are desired to make the imaged content appear correct and more readable. These approaches often refer to the distortion as *skew* and the removal process as *deskewing*.

From a single image it is difficult to correct distortion for documents with *arbitrary* shape. One interesting benefit of camera-based imaging is that 3D points on the document surface can be reconstructed by coupling an active lighting device with the camera. From these 3D points a triangulated mesh can be constructed that approximates the document's 3D surface [2, 16]. The problem now, as shown in figure 1, is given an image of a distorted document and a 3D reconstruction of the document's surface, can we produce a *restored* image of the document as it would appear in its planar format.

Paper documents can be modelled mathematically as *applicable surfaces* which are developable to a plane, i.e. the surface can be unfolded onto a plane without stretching or tearing. Thus, an exact 3D model of the document's surface can be mapped back to a plane without any distortion. In practice, however, this is not possible. Pilu [16] showed that the acquired 3D surface cannot be flattened by simple 2D re-tiling. Because the 3D mesh represents only an approximation of the true surface, subject to reconstruction errors and under-sampling, the reconstructed surface is not truly iso-metric 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 (i.e. distortion) must be allowed in the 3D to 2D mapping process.

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 [12], 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 angle distortion. In this paper, we show how conformal mapping can be used to parameterize the document's 3D surface to deskew an image of the warped document. Our experiments show that the conformal deskewing approach can restore images of

Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)

distorted documents to within a single pixel of their "flattened" representation. In addition, the conformal parameterization can be computed quickly by solving a linear system of equations and can incorporate knowledge about 3D surface points in the mapping process to establish orientation.

The remainder of this paper is organized as follows: section 2 discusses related work; section 3 discusses the deskewing algorithm; section 4 demonstrates the approach on several test cases; section 5 concludes the work.

# 2 Related Work

We discuss related work in two contexts: document restoration and 3D mesh parameterization.

## 2.1 Document Restoration

Distortion correction algorithms have traditionally focused on planar skew found in flatbed imaged items where the imaged content is not in alignment with the image axis. These approaches uses various techniques to compute affine and planar transforms to deskew the image (e.g. see [1, 11]). Another type of distortion commonly addressed is the curvature distortion near the spine of a book. While this distortion is caused by the material's non-planar shape, the restrictive and predictable nature of this deformation makes it relatively easy to correct using cylindrical models [4, 5, 19, 20].

Brown et al. [2] and Pilu [16] addressed *arbitrarily* distorted documents using the document's 3D shape. Structured-light devices are used with the imaging camera to reconstruct an approximation of the document's surface. The 3D surface model is flatten to a plane using relaxation methods that minimize spring energies between 3D vertices. The $z$ values of the document's 3D surface are gradually reduced to the $x$-$y$ plane. When the distance between adjacent vertices change, spring energies are induced. The relaxation method iteratively adjust vertex positions on the 2D plane until spring energies are minimized. A corrective warp between the 2D points in the original input image and their corresponding flattened vertices is used to create the restored image.

While these approaches are successful, they suffer from being computationally slow. Brown et al. [2] state that their approach takes roughly one minute to converge on a regularly sampled mesh with $45 \times 45$ vertices. Pilu states [16] that over 1000 iterations on a grid of $20 \times 15$ 3D points are needed. Moreover, the orientation of the resulting flatten representation cannot be predicted and an additional corrective transformation is needed to bring the flattened representation into a desired alignment.

## 2.2 Mesh Parameterization

Other related work can be found in the areas of computer graphics and geometric modelling, where 2D parameterization of 3D meshes is of great interest for remeshing, surface fitting, and to compute texture-coordinates. Approaches in this area vary by the distortion metric used and the underlying approaches employed.

Several approaches use spring-like energies between vertices [3, 10, 18] to minimize distance, angle distortion, or a combination of both. These differ from those used for document restoration in that they can solve the parameterization with a linear system, but require the boundary to be fixed. Such boundary restrictions is not possible for our problem since the 3D reconstruction obtained from the document is only of a sub-section of the document and the boundary is unknown.

Parameterization using conformal mapping has been addressed using Discrete Harmonic Maps [9], Differential geometry [17], and finite-elements approaches [13]. These approaches also require that boundary conditions be specified. Recently, Desburn [7] and Levy [15] presented techniques that can handle arbitrary boundaries. While these approaches typical require user interaction to help cut the closed 3D meshes into surface patches and are focused on parameterization for texture-mapping, the basic idea can be used by our restoration problem.

In the following section, we adapt the conformal mapping procedure for use in a document restoration framework. Our overall procedure is detailed.

# 3 Conformal Mapping

A mapping $f$ is *conformal at point* $a \in \Omega$ if for each pair of smooth curves $\nu$ and $\delta$ in $\Omega$ that pass through point $a$, the angle between their tangent vectors $\nu'(a)$ and $\delta'(a)$ at point $a$ is the same as the angle between the tangent vectors $(f \circ \nu)'(a)$ and $(f \circ \delta)'(a)$, i.e. the images of $\nu$ and $\delta$ under $f$. If $f$ is conformal *for all points* in $\Omega$ it is called a conformal mapping.

Consider now a continuous closed 3D surface represented by a vector $r(u, v)$, parameterized by $u$ and $v$, with components or $x_1, x_2, x_3$. The mapping $R^3 \mapsto R^2$ of $r(u, v)$ to the $u$-$v$ plane is conformal [6] if:

$$\nabla^2 r = \frac{\partial^2 r}{\partial u^2} + \frac{\partial^2 r}{\partial v^2} = 0 \tag{1}$$

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

For our application, since the input is a set of discrete 3D points on $r$ reconstructed on the document's surface, we are looking for an inverse mapping from $(x_1, x_2, x_3) \mapsto (u, v)$. Consider now a 2D function

$f(x, y)$, parameterized by $(x, y)$ that returns a 2D point $(u, v)$. The 2D Laplacian operator can be satisfied by the Cauchy-Riemann equation [14] such that:

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

which implies

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

The problem 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 $u$-$v$ plane in parameter space [15]. Each 3D triangle is converted to its local 2D coordinate frame aligned by the triangle's normal as shown in figure 2. Using this idea, we can use the reconstructed 3D points to form a mesh $M = \{\mathbf{p}_{i\in[1\ n]}, T_{j\in[1\ m]}\}$, where $\mathbf{p}_i$ are the $n$ 3D points acquired on the documents surfaces and serve as the vertices, and $T_j$ represents the resulting $m$ triangles, denoted as a tuple of three mesh vertices. In their local coordinate frame, we can consider $\mathbf{p}_i = (x'_i, y'_i)$ (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 original 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:

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

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

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

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

$$\begin{bmatrix} \frac{\partial u}{\partial x} - \frac{\partial v}{\partial y} \\ \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y} \end{bmatrix} =$$

$$\frac{1}{2A_T}\begin{bmatrix} \Delta x_1 & \Delta x_2 & \Delta x_3 & -\Delta y_1 & -\Delta y_2 & -\Delta y_3 \\ \Delta y_1 & \Delta y_2 & \Delta y_3 & \Delta x_1 & \Delta x_2 & \Delta x_3 \end{bmatrix}\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ \dots \\ v_1 \\ v_2 \\ v_3 \end{bmatrix}$$

$$= \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{4}$$

2D Triangle to Triangle Mapping

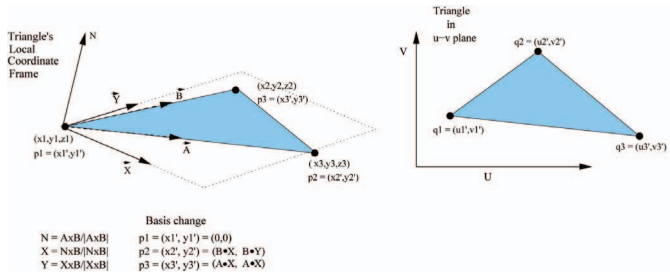


Figure 2: For each triangle, we convert its global $(x_i, y_i, z_i)$ representation to its local coordinate system. The system is defined with the triangles normal as the $z$-axis, so the triangle is embedded in the local $x$-$y$ plane. We now consider the mapping of the 2D $x$-$y$ local coordinates to the corresponding triangle in the $u$-$v$ plane.

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 equation 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, we can write a global system of equations, $\mathbf{Ax} = \mathbf{b}$, that incorporates all the vertices and triangles in mesh $M$, such that matrix $\mathbf{A}$ and vector $\mathbf{x}$ are defined as:

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

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 in equation 4. This results in $\mathbf{A}$ being a $2m \times 2n$ matrix. The entries for vector $\mathbf{x}$ are the desired conformal points $(u_i, v_i)$.

## 3.1 Solving the system

To obtain a unique solution for the system $\mathbf{Ax} = \mathbf{b}$ up to a similitude, 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, $\mathbf{p}_k$, to map to specified conformal locations $\mathbf{q}_k$, we modify


1063-6919/05 $20.00 © 2005 IEEE

the matrix $\mathbf{A}$ and vector $\mathbf{b}$ to reflect this constraint, such that $\mathbf{A}$'s $l$ columns entries, $a_k$ and $a_{2k}$, are removed from the matrix $\mathbf{A}$. A new $\mathbf{b}$ is constructed to incorporate this constraint in the solution, such that:

$$\mathbf{b} = -\begin{bmatrix} a_{k_1} & a_{2k_1} & \dots & a_{k_l} & a_{2k_l} \end{bmatrix} \begin{bmatrix} u_{k1} \\ \vdots \\ u_{k_l} \\ v_{k_1} \\ \vdots \\ v_{k_l} \end{bmatrix},$$

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

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

### 3.2 Deskewing Procedure

Our input consists of a triangulated 3D mesh of the non-planar document's surface and a 2D image of that surface. The 3D mesh is typically acquired using a structured-light scanner [2, 16]. In these approaches, the imaging camera is used together with an active lighting device (e.g. a light projector) to reconstruct the 3D surface. As a result, the location of each reconstructed 3D point, $\mathbf{p}_i$, on the 3D mesh has an associated 2D location in the camera's image plane $(s_i, t_i)$. These 2D points correspond to the distorted image we want to correct.

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 image 2D points to the conformal parameterization, such that $(s_i, t_i) \rightarrow (u_i, v_i)$. The restoration procedure is now a matter of warping the 2D points between the distorted image to their locations in the restored image. In practice, this *non-planar deskewing* is realized by piecewise affine warps between triangles in the distorted images to their corresponding triangles in the restored image, as shown in figure 3.

## 4 Results

We demonstrate our approach on printed patterns that have been arbitrarily distorted. A control image is available for comparison with our restored images. For each experiment, the document's 3D surface was reconstructed using a structured-light scanner with an estimated reconstruction accuracy of $0.31mm$, over a surface of approximately $20cm \times 15cm$. A 3D surface mesh is constructed from $46 \times 46$ 3D points, resulting in $2116$ vertices and $4050$ triangles. Corresponding images are captured using a standard NTSC camera, with $(640 \times 480)$ resolution. For more information on the 3D acquisition setup used for these experiments see [2].

Restoration Procedure

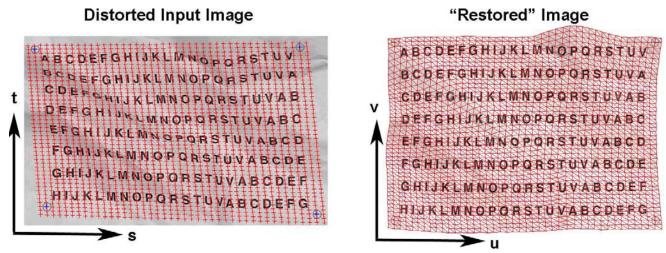


Figure 3: The restoration procedure is performed by piecewise warping 2D points in the original distorted image to the locations defined by the conformal map.

The experiments are performed on a Pentium III 1Ghz machine with 512MB. The procedure is implemented in Matlab and computes the conformal mapping and corrected image a matter of seconds.

### 4.1 Pixel Distances and Surface Area Change

Figure 4 shows five test cases. A document was imaged before it was warped (i.e. planar) and serves as the experiment's control. The five documents are restored using the conformal mapping procedure described in the previous section. Since illumination artifacts are not corrected, we cannot make a pixel-wise comparison, such as PSNR, between the restored images to the experimental control. Instead, we compare the distances between the *corners* of the checkerboard pattern in the *restored* and controlled images. There are $96$ corners in total. A homography is used to align the restored document's and control's extracted point[1] corners. Our experiments show that the corner features in the restored images are within a pixel of the control patterns checkerboard corners.

We also report the similarity (ratio) between the document's 3D surface area and the conformal map's area. Scale is normalized by scaling the conformal maps edges to a fixed edge length in the input mesh. Our results show

[1] While the orientation of the mapping can be controlled by fixing vertices, the conformal map cannot account for the perspective projection of the control image and a homography in needed to align the two patterns. In practice, this additional homography is not used.

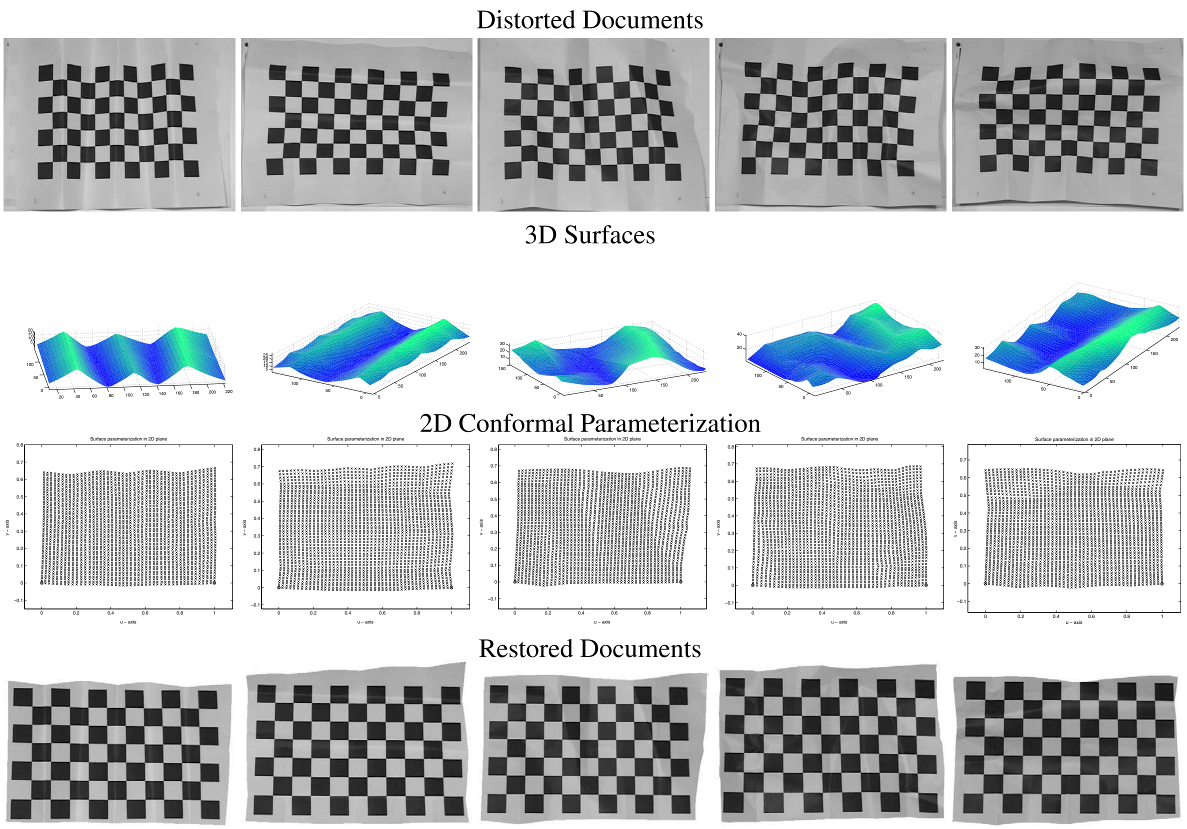


Control Image and Quantitative Results

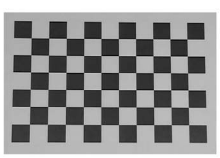

| | Testcase 1 | Testcase 2 | Testcase 3 | Testcase 4 | Testcase 5 |
|---|---|---|---|---|---|
| Mean Pixel Distance | 0.71 | 0.79 | 0.45 | 1.00 | 0.40 |
| Standard Deviation | 0.44 | 0.48 | 0.271 | 0.62 | 0.21 |
| | | | | | |
| Surface area similarity (between 3D and 2D mesh) | .987 | .988 | .991 | .989 | .982 |

Figure 4: Results of our restoration algorithm applied to five test cases. The figure shows the imaged documents and their corresponding 3D shape as a triangulated mesh. 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. Also the ratio of the surface areas between the 3D mesh and its flattened representation are given.

Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)
1063-6919/05 $20.00 © 2005 **IEEE**


Figure 5: Example using a lettered document. The top row shows the original and restored image. The bottom row shows the control image and the restored image thresholded for clarity.

that the overall surface area change between the input and restored area is very small, less than $2\%$ for all testcases.

### 4.2 Lettered Document

The distorted lettered document presented in figure 1 of this paper has been corrected in figure 5. In figure 5(top) shading effects have not been removed and may give the impression that the document is still warped, closer observation, however, will show that the document is restored. Figure 5(bot) row compares the restored image with a undistorted control image. The restored image was converted to a bi-tonal representation for clarity. The bi-tonal images shows the document's content has been corrected.

### 4.3 Vertex Placement

Figure 6 shows how prior knowledge of the document surface can be used to control orientation. Three different vertex selections have been used to constrain the conformal map as discussed in section 3.1. The first example has no prior knowledge about the 3D points and two corner points on the 3D mesh are chosen to mapped to conformal locations $[0, 0]$ and $[1, 1]$. In the second example, we constrain the points corresponding to the bottom checkerboard corners to map to $[0, 0]$ and $[1, 0]$, resulting in a horizontal orientation. In the third example, we constrain three vertices, the top corners and the lower left corner. The aspect ratio of the checkerboard pattern is known (11 horizontal blocks and 7 vertical blocks, i.e. $7/11$ or $0.63$), so we set the corresponding vertices to $[0, 0]$, $[0, 1]$ and $[-.63, 0]$, which results in a rotated orientation. Incorporating constraints for controlling orientation is another benefit offered by conformal mapping over the existing techniques. Previous work by [2, 16] required an additional user-specified transformation to bring the restored image into a desired alignment.

## 5 Discussion

Our results 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 relaxation algorithms; however, the conformal map can be computed in a matter of seconds. Furthermore, the conformal map technique allows us to incorporate prior knowledge into the mapping process to control orientation.

While 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 be need to be captured and registered to fill in missing intensity information. This is an interesting issue for further research. We also note that shading artifacts are not addressed. The presence of these in the restored image can give a strong impression of distortion. While this can be lessened with intensity adjustments, it should be possible to remove the shading artifacts given the approximation of the document's surface and is a subject for future work. Initial results on shading by [19] and [21] are promising, but are not yet applicable to arbitrary surfaces.

## 6 Conclusion

We have presented a novel technique to deskew non-planar documents using conformal mapping. By parameterizing the 3D surface to a 2D plane using the conformal constraint, we can compute a corrective map between the original image and a *restored* representation. This approach is significantly faster than previous algorithms, allows vertices to be constrained to control orientation, and can restore arbitrarily distorted document's to within a single pixel of their true planar format in image-space.

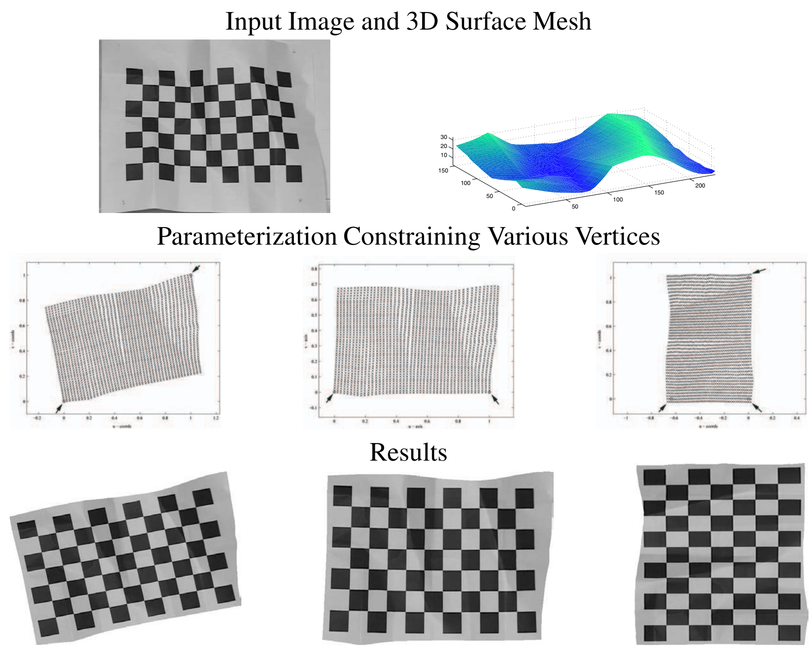


Figure 6: Example of different constraints on the vertices. Orientation can be controlled by specifying where certain 3D points show map to in the $u$-$v$ plane. Arrows point to the constrained vertices.

# References

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

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

[3] S. Campagna and H.-P. Seidel. Parameterizing meshes with aribitrary topology. In *Image and Multidimensional Digital Signal Processing*, pages 287–290, 1998.

[4] H. Cao, Ding, X., 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.

[5] H. Cao, X. Ding, and C. Liu. A clindrical surface model to rectify the bound document image. In *International Conference on Computer Vision, ICCV'2003*, pages 228–233, 2003.

[6] R. Courant. *Dirichelt's Principle, Conformal Mapping, and Minimal Surfaces*, chapter Chapter III: Plateau's Problem, pages 95–134. Springer-Verlag, 1977.

[7] M. Desburn, M. Meyer, and P. Alliez. Intrinsic parameterizations of surface meshes. In *EUROGRAPHICS 2002*, volume 21, pages 209–218, Saarbrucken, Germany, September 2002.

[8] D. Doermann, J. Liang, and H. Li. Progress in camera-based document image analysis. In *International Conference on Document Analysis and Recognition (ICDAR'03)*, pages p. 606–616, August 2003.

[9] M. Eck, T. Derose, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *SIGGRAPH'95*, pages 173–198, 1995.

[10] M. Floater and M. Reimers. Meshless parameterization and surface reconstruction. *Computer Aided Design*, 18(2):77–92, March 2001.

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

[12] A. Gray, editor. *Modern Differential Geometry of Curves and Surfaces*. CRC Press, second edition, 1998.

[13] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, H. Sapriro, and M. Halle. Conformal surface parameterization for texture mapping. *IEEE Transcations on Visualiation and Computer Graphics*, 6(2):181–189, April 2000.

[14] S.C. Krantz. *Handbook of Complex Analysis*, chapter The Cauchy-Riemann Equations. Birkhauser, Boston, MA, 1999.

[15] B. Lévy, S. Petijean, R. Nicholas, and J. Maillot. Least squares conformal maps for automatic texture atlas generation. In *ACM Transactions of Graphics(Proceedings SIGGRAPH 02)*, San Antonio, TX.

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

[17] U. Pinkall and K. Polthier. Computing discrete minimal surfaces. *Experimental Math*, 2(1):15–36, 1993.

[18] A. Praun, E. Finkelstein and H. Hoppe. Lapped textures. In *ACM Transactions of Graphics(Proceedings SIGGRAPH 2000)*, pages 465–470, New Orleans, LA, 2000.

[19] Y.C. Tsoi and M.S. Brown. Geometric and shading correction of imaged printed materials: A unified approach using boundary. In *IEEE Computer Vision and Pattern Recognition*, June 2004.

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

[21] Z. Zhang, C.L. Tan, and L. Fan. Estimation of 3d shape of warped documents surface for image restoration. In *IEEE Computer Vision and Pattern Recognition*, 2004.


1063-6919/05 $20.00 © 2005 **IEEE**