

Robust Segmentation of Historical Parchment XMT Images for Virtual Unrolling

Chang Liu

School of Astronautics
Beihang University
Beijing, China

School of Computer Science
Cardiff University, UK
LiuC42@cardiff.ac.uk

Paul L. Rosin

School of Computer Science
Cardiff University
Cardiff, UK

RosinPL@cardiff.ac.uk

Yu-Kun Lai

School of Computer Science
Cardiff University
Cardiff, UK

yukun.lai@cs.cardiff.ac.uk

Weiduo Hu

School of Astronautics
Beihang University
Beijing, China

weiduo.hu@buaa.edu.cn

Abstract—Historical parchment scrolls are fragile, and prone to damage from a variety of causes such as fire, water, and general mistreatment. Consequently many of these scrolls cannot be unrolled, so that their contents have remained hidden for centuries. To overcome these difficulties, we have developed a method of segmenting X-ray tomographic scans of parchment which enables a “virtual unrolling” of these documents. After an initial segmentation we link the broken layers of the parchment. Then, junction sections are extracted from the boundaries of the parchment. Subsequently, we find the fused regions which are formed by layers stuck together, and separate them into several layers by reconstructing the missing boundaries using parallel connecting curves. Experiments on the fifteenth century Bressingham scroll validate the effectiveness of our segmentation method.

Keywords—X-ray, parchment, parallel curve, flatten.

I. INTRODUCTION

Much of the valuable historical information of the western world is recorded on parchment, which is made from limed animal membrane that has been processed to be smooth and flat [1]. However, over hundreds of years, the layers of the scrolled parchments have become brittle and may also get stuck together, as demonstrated in Fig. 1. The parchment scroll in Fig. 1 is the Bressingham scroll, which is an account from the manor of Bressingham, dated 1408-9 (NRO, PHI 468/5) [1]. The records on the scroll include: the income of the lord from the manor and his expenditure, profits from holding the manor court, sales of underwood, and leasing out the fishing rights. The width of the scroll measures around 270 mm. The total length of the scroll is unknown, as it is impossible to unroll completely: at approximately 100 mm from the start of the parchment the scroll has become fused together. The fusing of the scroll is most likely to have been caused by exposure to moisture and damp storage. Any attempt to physically unfold the parchment document will lead to delamination of the surface of the parchment causing an unacceptable level of damage.

Digital document restoration technology initiates new methods for recovering and reconstructing those documents and information which cannot be accessed physically. This technology has been an extremely active area of research in recent years. So far, much attention has been paid to

the analysis of regular photographic images of historical documents and three-dimensional reconstruction and virtual flattening of deformed but non-scrolled parchment documents. 3D reconstructions of documents can be grouped into three different classes [2]: single-image methods, which reconstruct the shape of a document based on their geometric and shading information [3], stereo-image methods, which restore the 3D surface of document by stereo images [4], and structured-light scanning methods, which calculate the 3D shape of a document with a structured-light scanner [5]. Based on the restored 3D surface, many surface parameterization methods [6], [7] were proposed to find a mapping from document surfaces to planar domains with minimum distortion, which tackle the problem of recovering the original flat shape of distorted documents. Current efforts have provided a new level of accessibility to many valuable literary works. However, such methods cannot cope with parchments which cannot be physically unrolled. Virtual unrolling is required for such documents, though only very few results have been reported.

Recovering the information written on a fragile parchment scroll is a difficult task since the characters are wrapped inside the parchment, thus cannot be photographed by normal cameras. Therefore non-interventive methods were researched based on X-ray microtomography (XMT) [8] scanning and virtual unrolling [1]. XMT is a miniaturised version of medical human body CT scanning, designed to look at specimens much smaller than the human body, with a resolution in the region of microns, and the X-ray absorption contrast is strong for iron gall ink, which is the most commonly used ink in Europe between the 12th and 19th centuries. However, no commercial or research X-ray tomography scanners were capable of providing the high quality of scan image of scroll. In recent years Queen Mary University of London (QMUL) developed high contrast, high dynamic range XMT scanners, which allow the imaging of the small differences in X-ray absorption between inks and the parchment they are deposited upon. As long as the ink contains some metal and the parchment is not very heavily limed, it can successfully visualise the ink and potentially recover text (Fig. 2), which makes it possible to digitise and decipher the information in scrolled parchments [1] [9] [10].

The physically fused region, i.e. where there is no actual separation between the layers as shown in Fig. 2, makes the processing of parchment scrolls more challenging than

traditional deformed documents. Therefore the separation of parchment layers is the major problem for parchment scroll restoration [11] [12]. Some widely used image segmentation methods, such as snakes [13] [14] and max-flow-based graph cut [15] [16], are not effective for this problem because while they can effectively differentiate between the foreground and the background, they cannot recognise fused regions and separate them into parchment layers.

Samko et al. proposed a novel graph cut method with a shape prior in order to extract the parchment layers from the image. They then recreated the missing boundaries in the fused region from the boundary of the opposite side of the same layer or the closest preserved boundary [11] [12]. The shape-prior-based graph cut will thin the parchment layers, thus providing much fewer inter-layer connections between two layers. However, if the parchment layer is relatively thin (Fig. 2), the graph cut is likely to fragment the layers, which may have been complete, into many small parts, so this method cannot cope with the Bressingham scroll.

Mocella, et al. reported a method to reveal various letters hidden inside the carbonised papyrus scrolls, which were buried by the eruption of the Mount Vesuvius in 79 AD and discovered 260 years ago at Herulaneum, using X-ray phase-contrast tomography [17]. Because the ink on the scrolls never penetrated into the fibres of the papyrus, but sat on top of them, X-rays passing through the ink sections are slightly distorted, or slowed down relative to the X-rays passing through the papyrus. Measuring “phase contrast” can produce very detailed information of its internal structure. This method is able to reveal several Greek letters hidden in the mangled and charred scroll, but cannot restore the content written in the scroll.

In this paper, we present a new method to address the separation of parchment layers for virtual unrolling of parchments. The parchment image is first segmented using the Otsu thresholding algorithm [18], and then the foreground is processed by three main steps: layer connection, refinement of segmentation, and skeleton connection. The first step connects the broken layers of the parchment by a strip with the width of parchment layer. At the next step, we find the fused regions and separate them into parchment layers on the basis of the shape information of parchment slice. Sometimes, the parchment layer is so thin that it is likely to be cut off by our segmentation method. Thus in the last step, we extract the skeletons of the layers and link the skeletons which should belong to one layer. We give details of these steps in the following sections, followed by experimental results and discussions.

II. LAYER CONNECTION

Some areas of the Bressingham parchment have become scuffed and delaminated, so that in the X-ray slice the parts of the layers corresponding to those areas are missing. A typical example is shown in Fig. 3. In this section, we will link the broken parchment layers by a strip with the width of parchment layer. It can be seen from Fig. 3 that the two layers which should be linked together must have two endpoints which are very close to each other, and also around two such endpoints the two layers have similar trends. These two properties enable us to recognise the layers which should be linked. We first extract the skeleton of the parchment, and then eliminate all

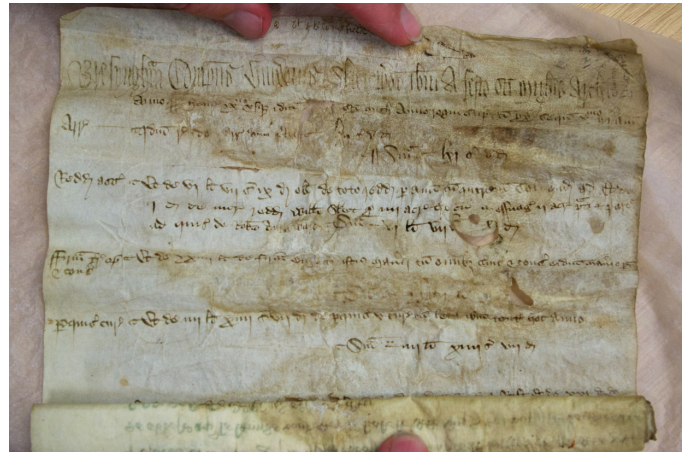


Figure 1. The Bressingham scroll, which cannot be physically unrolled.

spur points. Figure 4 depicts the skeletons of the two layers which should be merged. On account of the two properties as well as the fact that linking endpoints should not intersect with the existing parchment, two endpoints p_1 and p_2 are potentially linked if all of the following are satisfied: i) the direction vectors v_1 and v_2 respectively at p_1 and p_2 have an intersection angle larger than a threshold value (we set as 60°), ii) p_1 is close to p_2 , and iii) the line segment connecting p_1 and p_2 should not intersect with any existing parchment.

Accordingly, given the list of r endpoints, an $r \times r$ matrix M is used to record the cost of connecting a pair of endpoints, defined as the length of the line segment connecting these endpoints if they can be linked, and infinite otherwise. We use a greedy approach and repeat the following process until no further link can be made. We take the minimum element of M , assuming it is $M(u, v)$. If $M(u, v)$ is not infinite, we link the endpoints p_u and p_v by a line segment with the width of the parchment layer, and then eliminate the u th row, v th column, u th column and v th row of matrix M . The algorithm ends when the minimum element of M is infinite, which means that there are no remaining endpoints to be linked. The linking result of Fig. 3 is illustrated in Fig. 5.




Figure 2. One of the X-ray images showing a cross-section of the Bressingham scroll. The bright spots are ink.




Figure 3. The layers of the parchment are broken in many places.




Figure 5. The result of connecting the broken layers present in Fig. 3.




Figure 4. The description of two layers which should be linked.

III. REFINEMENT OF SEGMENTATION

In this part we aim to separate those fused regions into several parchment layers. The main steps of our segmentation refinement include junction point detection, fused region detection, parallel curve generation and missing boundary reconstruction.

A. Junction Point Detection

A junction point is a point which has large curvature, and we can use a purely geometric approach to detect junction points. Given a point p_0 on the boundary, we take n pixels $\{p_{-n}, p_{-n+1}, \dots, p_{-1}\}$ from the left neighborhood on the boundary of p_0 , and then take n pixels $\{p_1, p_2, \dots, p_n\}$ from its right neighborhood on the boundary (we set $n = 15$). If the intersection angle of vector p_0p_n and p_0p_{-n} is less than a threshold value (set as 120° in our algorithm), the point p_0 will be considered as a junction point. Figure 6 demonstrates the process of junction point detection and all detected junction points. It can be seen in this figure that some junction points are adjacent along the boundary, so that they form junction sections.

B. Fused Region Detection

A fused region is formed by some parchment layers merged together. As long as the fused regions are detected, the stuck layers can be immediately recovered by separating the fused regions. In fact, the junction sections are caused by fused regions, so a fused region must contain the middle point of at least a junction section. Given a binary image of a parchment slice (Fig. 7(a)), we connect the two endpoints of each junction section by a line, and then cut off along this line the region where the junction section is located. Now the binary image

has been separated into some small regions, as shown in Fig. 7(b). Consequently, the fused regions are the unions of those small regions which contain the middle points of junction sections (Fig. 7(c)).

C. Parallel Curve Generation

After detecting all fused regions, we will separate the fused regions by linking two junction sections with a curve parallel to the closest preserved boundaries to the two junction sections. In this part, we will describe a parallel curve generation method for linking two points by the curve which is most parallel to another curve.

Assume that c is an arbitrary curve and p_o, p_e are two arbitrary points in the plane, whose coordinates are (x_o, y_o) and (x_e, y_e) , as shown in Fig. 8. Our aim is to link p_o and p_e by a curve which is as parallel as possible to curve c .

First of all we approximate the curve c by an n -sided polygonal curve. The direction vector of each side of the polygonal




Figure 6. Some junction points adjacent along the boundary gather together and form a junction section.

curve is represented as (a_i, b_i) , $i = 1, 2, \dots, n$. Suppose that the curve connecting p_o and p_e can be approximated by an n -sided polygonal curve too, the direction vector of whose each side is denoted as (x_i, y_i) , $i=1,2,\dots,n$, as illustrated in Fig. 9. Therefore, the fact that the curve which links p_o and p_e is the most parallel to curve c is equivalent to the fact that the polygonal curve which links p_o and p_e is the most parallel to the polygonal curve which approximates curve c . Consequently, we obtain a cost function which has following form

$$\min \sum_{i=1}^n (x_i - a_i)^2 + (y_i - b_i)^2 \quad (1)$$

s.t.

$$\sum_{i=1}^n x_i = x_e - x_o, \sum_{i=1}^n y_i = y_e - y_o.$$

Equation 1 can be rewritten in matrix form

$$\min \mathbf{X}^T \mathbf{I} \mathbf{X} + \mathbf{P}^T \mathbf{X} + q \quad (2)$$

s.t.

$$\mathbf{A} \mathbf{X} = \mathbf{M}.$$

where, $\mathbf{P} = [-2a_1 - 2b_1 - 2a_2 - 2b_2 \dots - 2a_n - 2b_n]^T$, \mathbf{I} is the identity matrix, $\mathbf{X} = [x_1 \ y_1 \ x_2 \ y_2 \ \dots \ x_n \ y_n]^T$, $q = \sum_{i=1}^n (a_i^2 + b_i^2)$, $\mathbf{M} = \begin{bmatrix} x_e - x_o \\ y_e - y_o \end{bmatrix}$, and $\mathbf{A} = \begin{bmatrix} 1 & 0 & 1 & 0 & \dots \\ 0 & 1 & 0 & 1 & \dots \end{bmatrix}$. Using Lagrange multipliers, we can incorporate the constraints, and obtain the following system of equations

$$\begin{cases} \frac{\partial}{\partial \mathbf{X}} [\mathbf{X}^T \mathbf{I} \mathbf{X} + \mathbf{P}^T \mathbf{X} - \beta^T (\mathbf{A} \mathbf{X} - \mathbf{M})] = 0 \\ \frac{\partial}{\partial \beta} [\mathbf{X}^T \mathbf{I} \mathbf{X} + \mathbf{P}^T \mathbf{X} - \beta^T (\mathbf{A} \mathbf{X} - \mathbf{M})] = 0 \end{cases}, \quad (3)$$




Figure 7. Detection of fused regions. (a) The binary image of a parchment slice. (b) The separated binary image. (c) The detected fused region.




Figure 8. A arbitrary curve c and two arbitrary points p_o and p_e .




Figure 9. c and the curve connecting p_o and p_e are approximated by a n -sided polygonal curve.

where β is the vector of Lagrange multipliers. Equation 3 leads to the following form

$$\begin{cases} 2\mathbf{I} \mathbf{X} + \mathbf{P} - \mathbf{A}^T \beta = 0 \\ \mathbf{A} \mathbf{X} - \mathbf{M} = 0 \end{cases}. \quad (4)$$

As the result, we have

$$\begin{bmatrix} \mathbf{X} \\ \beta \end{bmatrix} = \begin{bmatrix} 2\mathbf{I} & -\mathbf{A}^T \\ -\mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{P} \\ \mathbf{M} \end{bmatrix}. \quad (5)$$

Expanding Eq. 5, we finally get

$$\mathbf{X} = \left(\frac{1}{2} \mathbf{I} + \frac{1}{2n} \mathbf{A}^T \mathbf{A} \right) \mathbf{P} + \frac{1}{n} \mathbf{A}^T \mathbf{M}. \quad (6)$$

Figure 10 shows a result of our linking method. Here the black curve is an arbitrary curve; the two green points are two arbitrary points, and the red curve is to connect the two green points and be as parallel as possible to the black curve.

D. Shape Energy Calculation

In this part we discuss a method of calculating shape energy of the parchment layers, which will serve the boundary




Figure 10. The red curve connects the two green points, while being as parallel as possible to the black curve.

reconstruction later. The shape energy of a point in a parchment layer reflects the closeness of the point to the layer boundary. The closer the point is to the layer boundary, the lower the shape energy of the point is. Given the parchment thickness m , we define shape energy on the transverse direction t of a layer to be a Gaussian function, as shown in Fig. 11.

On the condition that there are interlayer connections between some layers, the shape energy on the transverse direction t of the connected part of those layers should show the same trend as is demonstrated in Fig. 12.

Thus, given d_x the perpendicular distance of each pixel x in a layer to its closest boundary, the energy function has the following form

$$E(x) = e^{-\frac{(d_x - \frac{2k-1}{2}m)^2}{2\sigma^2}}, \quad (7)$$

where k is the parchment layer counter, $k = 1, 2, 3, \dots$, and d_x and m meet the following constraint

$$(k-1)m < d_x < km, \quad (8)$$

σ is the parameter which we estimate as $\sigma = m/3$, so that almost 99.7% of the energy of Gaussian function will lie within the layer. A result of shape energy calculation is shown in Fig. 13. It is evident that the shape energy reaches its peak in the middle of layer, and diminishes progressively from the middle of the layer towards the boundary.

E. Missing Boundary Reconstruction

In this part, we describe our method for recreating the missing boundaries. Provided there are k junction sections in the image, we initialise a $k \times k$ cost matrix \mathbf{W} with each element as infinite. Then we find the closest preserved boundaries for each junction section as follows. As depicted in Fig. 14, given a junction section R_i , whose two endpoints are l_i and r_i , a line which passes through l_i and r_i meets the left closest boundary at m_i and the right closest boundary at n_i , then these two closest boundaries on two sides of R_i are the closest preserved boundaries of R_i . Topologically, two junction




Figure 11. Shape energy on the transverse direction t of a layer is a Gaussian function.




Figure 12. The trend of the shape energy in the interlayer connections between several fused layers.




Figure 13. The shape energy only computed within the segmentation mask.

sections may be matched only if they are located on different boundaries but on the same fused region, as illustrated in Fig. 7. If two junction sections which are in different boundaries and in the same fused region have the same closest preserved boundary, we will check whether their middle points may be linked to reconstruct missing boundaries.

Providing that there exist two junction sections R_j and R_i , $i > j$, which are on different boundaries but in the same fused region, and have the same closest preserved boundary (Fig. 14), the two intersection points m_i and m_j respectively corresponding to R_i and R_j separate L into two parts, which are represented by the solid line and dot line respectively in Fig. 14. We only take into account the part which completely lies between the two lines $m_i l_i$ and $m_j l_j$, that is, the solid part in Fig. 14.

First of all, we use the solid line part to generate a curve Q which connects the middle points of R_i and R_j by the method in section III.C, and then check whether Q intersects the existing boundaries at any places other than R_i and R_j . If not, there is a possibility that the region between R_i and R_j is an interlayer connection, so we calculate the energy between




Figure 14. Two junction sections on different boundaries but in the same fused region have a same closest preserved boundary.

R_i and R_j along curve Q by the following form.

$$H = \frac{1}{g} \sum_{s=1}^g E(x_s) + \lambda \left[1 - \frac{\min(m_i l_i, m_j l_j)}{\max(m_i l_i, m_j l_j)} \right], \quad (9)$$

where x_s denotes the pixel of Q , g is the number of pixels in Q , and E is the shape energy calculated by Eq. 7. In Eq. 9, the first term $\frac{1}{g} \sum_{s=1}^g E(x_s)$ measures the mean shape energy of Q , and the second term $1 - \frac{\min(m_i l_i, m_j l_j)}{\max(m_i l_i, m_j l_j)}$ reflects the similarity of the distances from R_i and R_j to their closest preserved boundary. The more similar the distances from R_i and R_j to their closest preserved boundary, the less this term. λ is a parameter which specifies the significance between these two terms. If $H < \mathbf{W}(i, j)$, we will set $\mathbf{W}(i, j) = H$.

After we apply the above method to all possible pairs of junction sections, if the minimum element of \mathbf{W} is not infinite, we will separate the fused region along the parallel curve linking the middle points of R_u and R_v , which corresponds to the minimum element $\mathbf{W}(u, v)$ of \mathbf{W} , $u > v$. Subsequently we eliminate the u th row, v th column, u th column and v th row of matrix \mathbf{W} , then update the already existing fused regions and let the algorithm begin all over again. The algorithm will terminate when the minimum element of \mathbf{W} is infinite, which means that there are no junction sections to be matched, or there is a single boundary in the image. Figure 15 illustrates a segmentation result of our algorithm.

IV. SKELETON CONNECTION

After finishing segmentation, we need to extract the skeleton of the parchment layer for virtual unrolling. However, sometimes some parts of the parchment layer are too thin to provide sufficient single pixel edges for our segmentation algorithm, then our segmentation algorithm will eliminate such very thin parts to guarantee the junction section detection and missing boundary reconstruction. The skeletons of these broken layers will mislead the virtual unrolling method to generate a wrong result. Thus before virtual unrolling, we need to link the skeletons of the layers which are mistakenly broken. The linkage method is similar to the method of missing boundary reconstruction in section III.E. Figure 16 illustrates the effect of our skeleton connection.

V. EXPERIMENTS

We demonstrate the performance of our segmentation method to real parchments, which vary in size and complexity. The data were acquired through tomographic development in the School of Medicine and Dentistry at QMUL [11]. We test our algorithm on a desktop PC with a 2.27GHz processor.

The first scroll is a small relatively simple example ($430 \times 430 \times 760$ pixels) to test our algorithm. The largest fused region in this parchment consists of three layers, and each



Figure 15. The result of our segmentation for parchment slice. (a) A fused region in the original image. (b) The segmentation of the fused region.




Figure 16. Skeleton connection. (a) The segmented parchment slice. It can be seen that a layer is mistakenly broken. (b) The skeleton image of the broken layer. (c) The linked skeleton provided by our algorithm.

layer is about 14 pixels wide ($m = 14$). It is in fairly good condition because it has relatively thick and even layers and the layers are complete. Accordingly we set $\lambda = 0$ in this test. Figure 17 illustrates the stages of our segmentation method. It can be seen that all the fused regions are correctly separated into several layers. Because the segmentation results for all the slices are quite similar, the example in Fig. 17 represents the performance of our algorithm for the whole data set.

In order to verify the correctness of our segmentation method, we extract the skeleton line from each of our segmentation results, and then use them to flatten the parchment scroll by the method [12] to recover the text written on the parchment. The flattening result is illustrated in Fig. 18. There are visible letters on the virtually unrolled parchment, demonstrating the effectiveness of our method.

The next parchment is the Bressingham scroll, which is




Figure 17. The result of our segmentation method to a parchment. (a) The original image. (b) The foreground extracted by Otsu's method. (c) The segmented parchment by our method. (d) A close-up view of a fragment. (e) Skeleton image.

much more complicated than the first small one. The part of the parchment being processed in this work consists of 1500 frames, with 1256×816 pixels in each frame. The largest fused region is comprised of more than six layers. In addition, not only are the layers uneven, but also they are split into many parts. The average thickness of the layers is only about six pixels. Furthermore, there exist some artifacts in some frames, which probably resulted from the laced plastazote enclosure which was used to safeguard the scroll during transportation from the archive to QMUL and during the scanning processes [1]. All of these pose a serious challenge to the segmentation for this parchment. Because the layers are too thin, we resize the image to double the original width and height before processing. We set $m = 11$, and $\lambda = 1$ for this data. Figure 19 demonstrates an example of the segmentation of Bressingham scroll. Figure 19 (a) shows the original image. The foreground extracted from Fig. 19 (a) is shown in Fig. 19 (b). It can be seen that there exists some artifact above the parchment and the parchment layer is broken near the centre. First of all we link the broken layer together, as shown in Fig. 19 (c), then segment the slice by our segmentation method (Fig. 19 (d)). Figure 19 (e) illustrates some fragments of Fig. 19 (d), which indicates that the fused regions are correctly divided into several layers. Figure 19 (f) shows the skeleton of the parchment extracted from Fig. 19 (d). In this figure, the skeleton is complete and no noise or spur points exist. To sum up, it can be concluded from Fig. 19 that our algorithm is effective to deal with the complicated parchment.

Figure 20 and Fig. 21 illustrate the Bressingham scroll virtually flattened on the basis of our segmentation result. Because the unrolled parchment scroll is so long, we just show some typical parts of the whole scroll, which can be representative of the whole one. In Fig. 20, the right image is the photo of the visible part of the Bressingham scroll, and the left image is the recovery of the same section by our algorithm. It is clear from comparison that the characters deciphered from the parchment are correct and legible enough, although the parchment has many layers broken and stuck together. Thus Fig. 20 and Fig. 21 are strong evidences that our method can correctly segment the parchment scroll for virtual unrolling.

VI. CONCLUSION

In this paper we present a novel method for segmenting the X-ray images of parchment scrolls for virtual unrolling. First, we connect the layers of the parchment which are broken. Second, the junction sections are detected from the edge map of the parchment. Subsequently, the fused regions are found and separated into several layers by means of the missing boundary reconstruction and parallel curve connection. The experiments indicate that our method is capable of segmenting a challenging historical parchment – the Bressingham scroll – and can make the physically unopenable scroll readable.




Figure 18. The revealed text from the small parchment.




Figure 19. Segmentation for Bressingham scroll. (a) The original image. (b) The foreground extracted from the original image. (c) The broken layer is connected. (d) The segmented image. (e) Some scaled fragments. (f) Skeleton.

ACKNOWLEDGMENT

The scroll was provided by the Norfolk Record Office and scanned as part of EPSRC project EP/G010110/1, High definition X-ray microtomography and advanced visualisation techniques for information recovery from unopenable historical documents. We would like to thank David Mills and Graham Davis (Queen Mary University), Antoinette Curtis (NRO) and the China Scholarship Council (File No. 201406020068).

REFERENCES

- [1] D. Mills, A. Curtis, G. Davis, P. Rosin, and Y.-K. Lai, "Apocalypso: revealing the Bressingham roll," *Journal of Paper Conservation*, vol. 15, no. 3, pp. 14–19, 2014.
- [2] K. Pal, M. Terras, and T. Weyrich, "3D reconstruction for damaged documents: imaging of the great parchment book," in *Int. Workshop on Historical Document Imaging and Processing*, 2013, pp. 14–21.
- [3] T. Wada, H. Ukida, and T. Matsuyama, "Shape from shading with interreflections under a proximal light source: distortion-free copying of an unfolded book," *Int. Journal of Computer Vision*, vol. 24, no. 2, pp. 125–135, 1997.
- [4] H. I. Koo, J. Kim, and N. I. Cho, "Composition of a dewarped and enhanced document image from two view images," *IEEE Trans. on Image Processing*, vol. 18, no. 7, pp. 1551–1562, 2009.
- [5] M. S. Brown, M. Sun, R. Yang, L. Yun, and W. B. Seales, "Restoring 2D content from distorted documents," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 11, pp. 1904–1916, 2007.
- [6] K. Pal, C. Schüller, D. Panozzo, O. Sorkine-Hornung, and T. Weyrich, "Content-aware surface parameterization for interactive restoration of historical documents," *Computer Graphics Forum*, vol. 33, no. 2, pp. 401–409, 2014.
- [7] O. Sorkine and M. Botsch, "Interactive shape modeling and deformation," in *Proc. Eurographics Tutorials*, 2009, pp. 11–37.

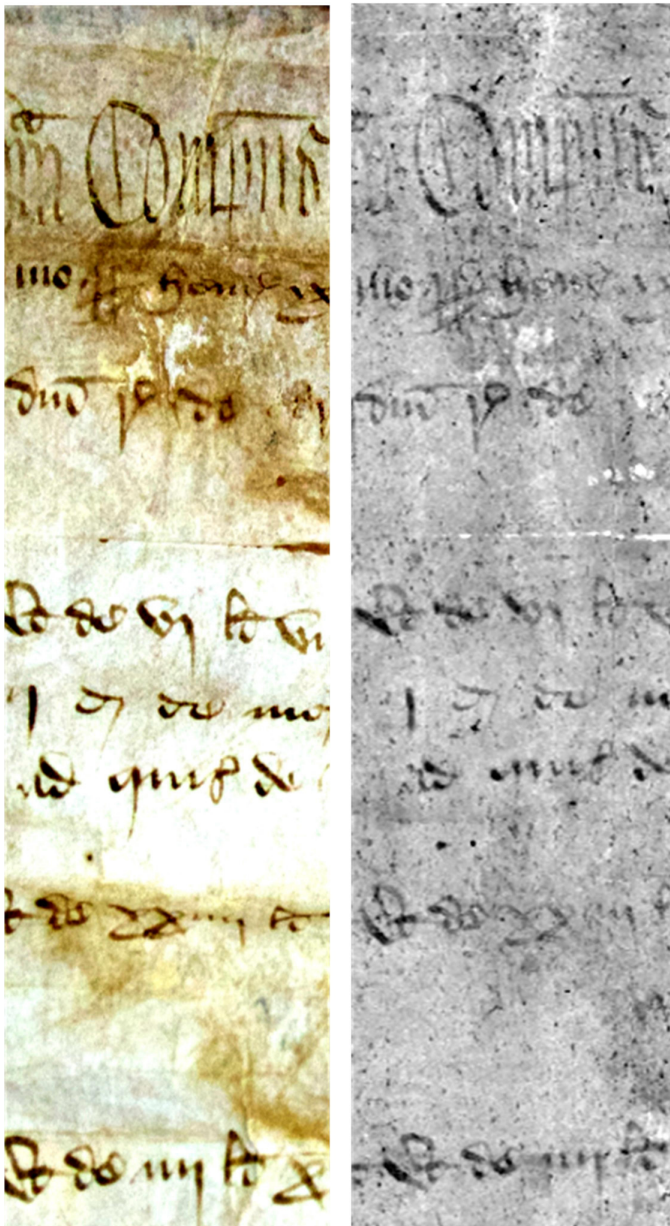


Figure 20. The visible part of the Bressingham scroll and its recovered image.

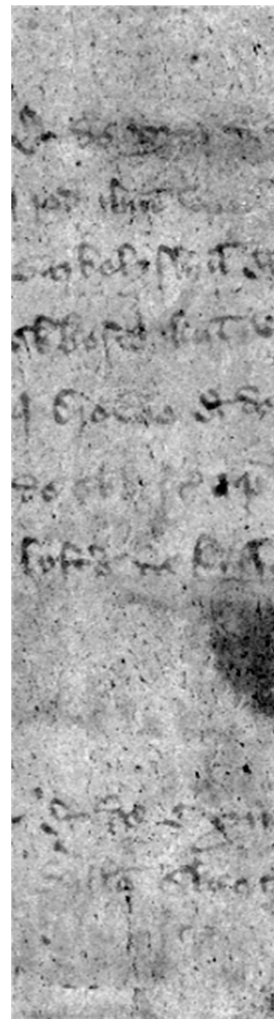


Figure 21. An unseen section of the Bressingham scroll (cannot be unrolled) which is visualised after virtual unrolling.

[8] G. Davis and J. Elliott, "High definition X-ray microtomography using a conventional impact X-ray source," in *Journal de Physique IV*, vol. 104, 2003, pp. 131–134.

[9] D. Mills, G. R. Davis, Y.-K. Lai, and P. Rosin, "Apocalypso: revealing lost text with XMT," in *SPIE Optical Engineering and Applications*, vol. 9212, 2014.

[10] F. Albertin, A. Astolfo, M. Stampanoni, E. Peccenini, Y. Hwu, F. Kaplan, and G. Margaritondo, "X-ray spectrometry and imaging for ancient administrative handwritten documents," *X-Ray Spectrometry*, vol. 44, no. 3, pp. 93–98, 2015.

[11] O. Samko, Y.-K. Lai, D. Marshall, and P. L. Rosin, "Segmentation of parchment scrolls for virtual unrolling," in *British Machine Vision Conference*, 2011, pp. 1–11.

[12] —, "Virtual unrolling and information recovery from scanned scrolled historical documents," *Pattern Recognition*, vol. 47, no. 1, pp. 248–259, 2014.

[13] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. Journal of Computer Vision*, vol. 1, no. 4, pp. 321–331,

1988.

[14] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow," *IEEE Trans. on Image Processing*, vol. 7, no. 3, pp. 359–369, 1998.

[15] Y. Y. Boykov and M.-P. Jolly, "Interactive graph cuts for optimal boundary and region segmentation of objects in ND images," in *Int. Conf. Computer Vision*, vol. 1, 2001, pp. 105–112.

[16] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124–1137, 2004.

[17] V. Mocella, E. Brun, C. Ferrero, and D. Delattre, "Revealing letters in rolled Herculaneum papyri by X-ray phase-contrast imaging," *Nature Communications*, vol. 6, 2015.

[18] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. on Systems, Man and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.