Segmentation of Parchment Scrolls for Virtual Unrolling

Oksana Samko
O.Samko@cs.cardiff.ac.uk
Yu-Kun Lai
Yukun.Lai@cs.cardiff.ac.uk
David Marshall
Dave.Marshall@cs.cardiff.ac.uk
Paul L. Rosin
Paul.Rosin@cs.cardiff.ac.uk

School of Computer Science & Informatics
Cardiff University
Cardiff, UK

There is a critical need to access the valuable information in historical scrolls that cannot be read by conventional means. In some cases, their physical deterioration is at such an advanced state that any attempt to unravel the document manually would cause catastrophic fragmentation, destroying the internal information. Use of X-ray microtomography, a new direction in digital document analysis, provides a digital copy of a scrolled parchment as a 3D volumetric object, see Fig. 1. Clearly, the parchment layers need to be separated to perform virtual unrolling and apply further digital document restoration methods.

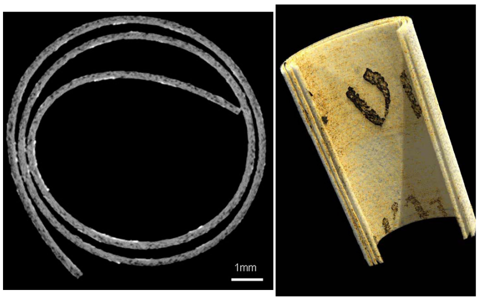

Figure 1: A small cut sample from a historical parchment scanned with the high definition XMT scanner. Left: cross-sectional tomographic slice; right: volume rendered cutaway view with pseudo-coloring.

Parchment is essentially animal skin and therefore has an irregular sponge-like structure, also its thickness may vary across a document surface. As a result of degradation over time, parchment may convert to its entropic form, gelatin, making the boundary between its layers difficult to observe even with the human eye. Image noise, low contrast and scanning artifacts may lead to even more indistinct parchment structure boundaries. A general algorithm can destroy damaged areas because of parchment's latent texture (oversegment), and not split tightly connected layers with zero gradient (undersegment) at the same time. The topology of the parchment smoothly changes from slice to slice, but can differ significantly across the whole scroll. The parchment ink thickness is only a few voxels deep (represented by the light pixels close to the parchment boundary), so it is very important to carefully process the boundary to avoid losing important information due to incorrect segmentation.

This paper describes the segmentation process utilised in our system for virtual unrolling of parchments, whose goal is to extract the scrolled parchment surface from the volumetric data, and to separate its touching layers. We exploit both geometric and pixel intensity features from the data. Our algorithm consists of three main steps: data filtering, segmentation, and postprocessing using geometric constrains. The first step, anisotropic filtering, makes the parchment structure more homogeneous, simultaneously preserving the parchment's layer boundaries. Using the CED filter [4] enables us to preserve the topology of the parchment layers, while the internal variation caused by the parchment's sponge-like structure is diminished. In the paper we discuss how to correctly adjust the CED parameters depending on the parchment size and condition.

At the next step we introduce the main segmentation routine, based on Graph Cut [1], with a novel shape prior optimisation. Here we incorporate parchment layer thickness information together with the traditional pixel intensity. Assume P is the set of image pixels, f_p is the binary label assigned to pixel p, N is the set of neighbourhood pixel pairs, N_u is the pairs of neigbours from the background and the dilated foreground. We define our energy as the weighted sum of data, smoothness and shape prior terms:

$$E(f)=\sum_{p\in P} D_p(f_p)+\lambda\sum_{(p,q)\in N} w_{pq}\varsigma(f_p,f_q)+\mu\sum_{(p,q)\in N_u} s_{pq}\varsigma(f_p,f_q). \tag{1}$$

Here $\varsigma(f_p,f_q)$ is 0 if $f_p=f_q$ and 1 otherwise, w_{pq} and s_{pq} are the smoothness and the shape weights (see the paper for full details), parameters λ and μ specify the relative importance between the terms.

This makes the segmentation more robust, however a few very tight connections between layers may still be retained. The reason for such incorrect connections is that the local boundary features may not exist, or it may be difficult to detect them reliably using the global optimisation. Therefore instead of involving time-consuming user interaction which also requires great accuracy from the user, in the last step of our algorithm we employ local geometric constraints to automatically separate such connections. The idea here is to recreate a missing boundary from the preserved boundary of the opposite side of the same layer, otherwise from the closest preserved boundary.

We evaluate our algorithm by applying it to segment three different parchment data sets, which vary in the parchment's condition, size and number of layers. Our experiments indicated that we were able to fully separate these parchment's layers, outperforming traditional segmentation methods such as Graph Cut and snakes both visually and numerically. To provide a numerical evaluation of our algorithm, we use multiple benchmark criteria: Probabilistic Rand Index, Variation of Information [3], as well as boundary-based precision-recall framework [2]. Fig. 2 shows the segmentation routine of our method applied to the two real parchment examples.

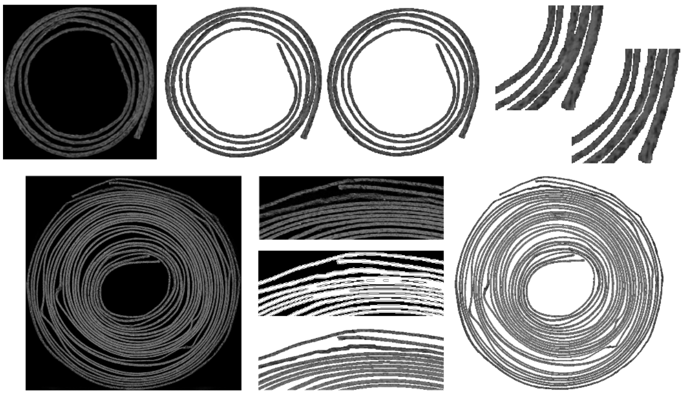

Figure 2: Top row: original slice; segmented by Graph Cut with shape prior; postprocessed slice; separated fragment before and after postprocessing. Bottom row: original slice; fragments: original, thresholded with assigned boundary, result; slice segmented by our method

[1] Y. Boykov and M.P. Jolly. Interactive graph cuts for optimal boundary and region segmentation. In *Proceedings of ICCV*, volume 1, pages 105–112, 2001.

[2] D.R. Martin, C.C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. PAMI*, 26:530–549, 2004.

[3] M. Meila. Comparing clusterings: an axiomatic view. In *Proc. of ICML*, pages 577–584, 2005.

[4] J. Weickert. Coherence-enhancing diffusion filtering. *Journal of Computer Vision*, 31:111–127, 1999.