Towards Automating Volumetric Segmentation for Virtual Unwrapping

Supporting Deep Learning Through Volumetric Page Instance Segmentation

Kristina GESSEL, University of Kentucky, United States
Stephen PARSONS, University of Kentucky, United States
Clifford PARKER, University of Kentucky, United States
William SEALES, University of Kentucky, United States

Keywords: *Volumetric Segmentation — X-Ray Microtomography — Virtual Unwrapping — Machine Learning*

CHNT Reference: Gessel K., Parsons S., Parker C. S., Seales W. B. 2020. Towards Automating Volumetric Segmentation for Virtual Unwrapping. Proceedings of the 25th International Conference on Cultural Heritage and New Technologies 2020, Vienna. DOI:xxxxxxx.

Introduction

A new algorithm for volumetric segmentation that produces high-quality segmentations of the pages of a manuscript is presented. This approach replaces the current segmentation component of the virtual unwrapping pipeline and greatly reduces the effort and time required of a user. The algorithm is applied to extract pages with clear ink signal from micro-CT scans of the M.910 manuscript. Future applications of this algorithm, particularly its ability to generate training data for a supervised machine learning algorithm that performs segmentation fully automatically, are also discussed.

Acquired in a severely damaged state by the Morgan Library & Museum in 1962, M.910 contains one of the earliest known full copies of Acts of the Apostles in any language. Unfortunately, however, only two exposed pages (Fig. 1 C) are readable. Fire and water damage charred and distorted the book's thin parchment pages, causing them to become gelatinized, brittle, and interlocked. Opening and turning them is impossible without causing more harm. In addition, moisture exposure corroded the iron-gall ink in some places, further deteriorating the parchment substrate.

The virtual unwrapping software pipeline can be used to noninvasively extract and recover the contents of damaged manuscripts too fragile to be opened, such as M.910. While any type of volumetric data can be used, in practice the manuscript is scanned with a micro-CT scanner to obtain tomographic data that shows the inner structures of the manuscript in thin, two-dimensional 'slices'. Multiple micro-CT scans of M.910 were completed in a 2017 scan session, with varying resolutions and energies.

The virtual unwrapping pipeline itself consists of three steps: first, segmentation aligns a mesh to a single page or layer in the tomographic data; next, the mesh is textured, or colored, by analyzing the intensities local to the mesh in the tomographic data; and finally, flattening removes folds, rolls, and warping in the page to produce a flattened page.

The existing segmentation algorithms in our virtual unwrapping software library, Volume Cartographer, create a bottleneck due to their user-driven nature. Initially, the user manually aligns a curve to the page or layer of interest in a single slice image (Fig. 1 B). An algorithm helps propagate the curve through subsequent slices using a variant of active contour modeling, but this proves challenging because the individual pages change shape, torn sections appear, and pages appear fused together so that it is difficult for even a human to distinguish one page from another. In practice, a user must manually intervene to correct and restart the segmentation after a small number of slices, typically less than 40 in the case of the M.910 manuscript. With such small batches, segmenting a single full page requires a dedicated user to spend more than one 8-hour work day. M.910 contains over 100 pages, so it is not feasible to complete the entire manuscript with this approach. A more automated approach to segmentation is necessary to eliminate the bottleneck.

Automating Segmentation

Automating the segmentation of volumetric data is a problem under active investigation in other areas, such as medical imaging. Many automated approaches favor machine learning techniques. For example, Januszewski et al. (2016) demonstrate the performance of their Flood-Filling Network on connectomics datasets.

Obtaining sufficient amounts of ground truth data to train a neural network to segment pages or layers of a manuscript is challenging. No public datasets with suitable labeled training data exist, so one must be generated. Manually labeling sufficient amounts of data to train a neural network would require a tremendous investment of time and manual labor, and in some cases the user cannot identify what the ground truth should be.

This problem motivated an approach to segmentation, based on computer vision concepts, that does not use machine learning itself but produces accurate segmentations--with only occasional assistance from a user--that can be repurposed as training data for a neural network.

As with the existing algorithm, the user provides a curve which traces the layer, and the algorithm attempts to propagate this curve to the next slice. The points that make up the curve are used to seed a modified flood fill algorithm, which produces a mask for every voxel in the layer (Fig 1 A). This mask is then thinned to create a new curve which more accurately follows the center of the layer (Fig. 1 B). The new curve is then used to seed the flood fill algorithm for the next slice, and the process repeats.

In some cases, the new algorithm will extend the mask of the page into neighboring pages--for example, if pages touch for prolonged periods of time. These errors can be manually removed after the segmentation process is complete, or the user can choose to immediately correct the segmentation when these errors occur. Correcting these errors is a relatively quick process.

Fig. 1: A: The mask generated by the algorithm, shown inside the tomographic data. B: The curve generated by thinning the mask. C: The two exposed pages of M.910.

Virtually Unwrapping M.910

Fig. 2 illustrates the results obtained from performing the virtual unwrapping process on the tomographic data using the new segmentation algorithm. The iron gall ink of the manuscript generally shows well in CT scans due to its iron content, but the readability of each page depends on the varying concentrations of ink within each letterform. This variation is visible in the exposed text shown above in Fig. 1 C, where the writing in the bottom right corner of the right page appears faded, especially when compared to the dark text on the left-hand page. The virtually unwrapped page shown in Fig. 2 C is partly legible but requires further enhancement to improve readability in the areas where ink signal is weak.

In addition, the pages of M.910 possess ink on both sides, a novel challenge for the virtual unwrapping pipeline that is compounded by the extraordinarily thin nature of M.910's parchment (0.03 mm). The fragment shown in Fig. 2 B nonetheless demonstrates that the new algorithm can help overcome this challenge by creating a more accurate segmentation that enables the virtual unwrapping pipeline to isolate ink on either side of the page.

As previously mentioned, the new approach to segmentation greatly increases a user's efficiency. A single page in M.910 was segmented using both the existing and the new segmentation methods to compare the quality of the resulting segmentations and the time each method takes to segment a page. With the existing algorithm, the page was fully segmented in 11 working hours. The same page, however, was fully segmented in only five hours using the new algorithm. The current approach requires the user's full attention for the entire segmentation process because it can successfully handle small batches of less than 40 slices in most cases. The new algorithm is capable of running independently for hundreds of slices in many cases, vastly reducing the amount of time a user spends on segmentation. For example, in the example above, the user adjusted the curve only 27 times through 4,928 slices. Adjusting the curve and resuming the segmentation typically takes less than five minutes, and in this example the user spent roughly two hours correcting and restarting the segmentation. Since the new algorithm is less reliant on user input, a single user can segment multiple pages concurrently.

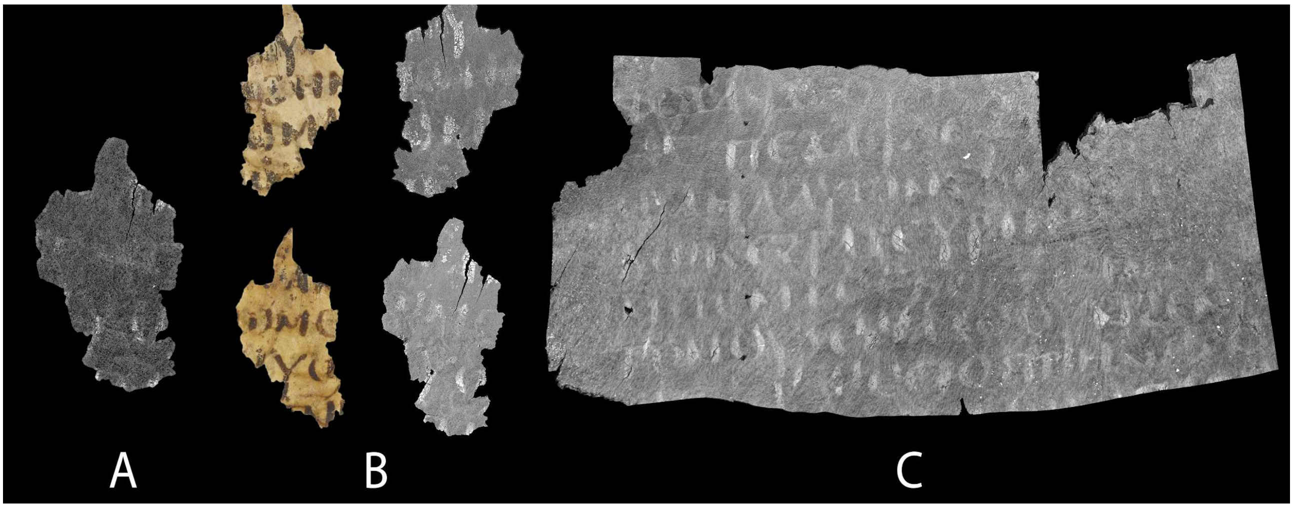

Fig. 2: A: The segmentation of a fragment of M.910. B: Ink on each side of this fragment is visible after the segmentation is processed with the virtual unwrapping pipeline. Color photographs of both sides of the fragment are included for comparison C: The segmentation of a page from the center of M.910, showing strong ink signals.

Conclusions

The new algorithm presented in this paper significantly reduces the amount of necessary input from a user while producing high-quality segmentations. The masks produced by this algorithm cover the page fully in most cases, so these masks could be used to train a supervised machine learning algorithm capable of segmenting the pages in a fully automated manner. In the particular case of M.910, the damaged text can be fully segmented in a semi-automated manner with the new algorithm. With additional ink enhancement strategies such as those addressed by Parker et al. (2019), pages that have been unreadable for hundreds of years may soon be readable again.

Acknowledgements

The authors would like to thank Dr. Paul Dilley from the University of Iowa for the study of revealed texts; Maria Fredericks, the Sherman Fairchild Head of the Thaw Conservation Center at the Morgan Library & Museum and her conservation team who made the scans possible and worked with us to ensure the manuscript's safety; MicroPhotonics for the use of the micro-CT equipment to scan the manuscript; Raj Manoharan, Seth Hogg, and Benjamin Ache of MicroPhotonics Inc. for micro-CT scanning services and consultation; and Christy Chapman for editorial support and input.

References

Parker C. S., Parsons S., Bandy J., Chapman C., Coppens F., Seales W. B. (2019). From invisibility to readability: Recovering the ink of Herculaneum. PLoS ONE 14(5): e0215775. https://doi.org/ 10.1371/journal.pone.0215775

Seales W. B., Parker C. S., Segal M., Tov E., Shor P., Porath Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances 2(9), p.e1601247.

Januszewski, M., Maitin-Shepard, J., Li, P., Kornfeld, J., Denk, W., Jain, V. (2016). Flood-Filling Networks. arXiv [Preprint]. Available at: https://arxiv.org/abs/1611.00421 (Accessed: 28 July 2020).