

Semi-Automated Segmentation of Brain MRI

600.446 Computer Integrated Surgery II

Final Project Report

Team members: Nathaniel Tippens, Alexander Liu

Mentors: Dr. Alfredo Quiñones-Hinojosa, Dr. Hadie Adams, Dr. Russell Taylor

Abstract

In this project, we developed an interactive watershed transform tool for the segmentation of MRI images of glioblastoma multiforme patients. This assisted segmentation tool increases accuracy and reduces inter and intra-observer variability present in current segmentation practices. A C++ implementation of this algorithm was developed within the Insight Toolkit (ITK) library, and further viscous modifications were made to the algorithm to improve the segmentation results of ambiguous, low-resolution scans. Finally, the algorithms were integrated into ITK-SNAP, a free open-source segmentation platform, in preparation for the algorithm's eventual application to volumetric studies.

Introduction

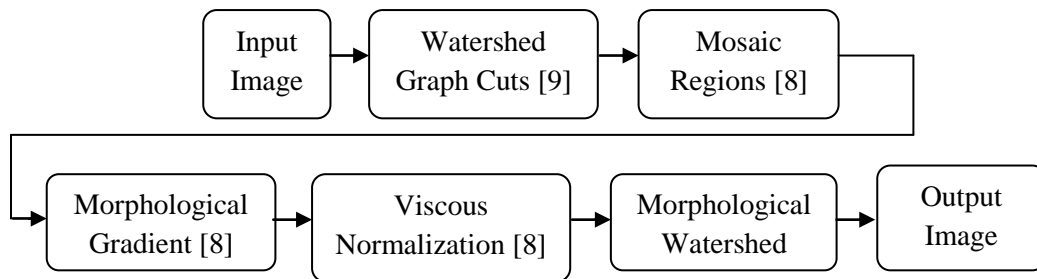
There is a growing amount of literature emphasizing the need for a standard method of medical image segmentation, due to the high variability and slowness of manual segmentation [1, 2]. Existing volumetric software tools are not practical for clinicians/researchers due to cost or complexity. In addition, the current golden standard of medical image segmentation requires the observer to manually delineate tumor contours on each slice of an MRI scan. This tedious process results in disagreement between observers, and it is often difficult for the same observer to reproduce segmentations. Moreover, neurosurgeons and radiotherapists, who heavily depend on these segmentations, have recently recognized the impact these variabilities can have on treatment outcomes. It has already been shown that the creation of software-assisted standards to define the final tumor contour can minimize operator variability, which may drastically improve patient outcomes [6]. The development of software to assist in tumor delineation would not only speed up segmentation time but also create more rigorous standards for the segmentation of target volumes, removing subjectivity from current practice and introducing a more quantified objective one. Apart from the clinical applications, such a tool is also invaluable to scientists

who wish to quantify tumor volumes accurately for research purposes. For the last decade, neurosurgeons have attempted to correlate extent of brain tumor resection with patient survival. However, these studies have shown heterogeneous results, with only a few studies attempting to answer this question with volumetric software, all of which employed manual segmentation to obtain volume [3, 4, 5, 6]. In the modern scientific era, an objective and reproducible segmentation method must be used as a new golden standard to mitigate the subjectivity inherent in current practice.

In this project, we developed an interactive watershed transform tool to fulfill the growing demand for medical volumetric software. This tool, which was a modification of ITK-SNAP [7], is freely available to researchers looking to quantify volumetric analysis of MRI tumor images in a more consistent, accurate manner.

Technical Approach

Initially, we planned to use the classical watershed algorithm to segment the image via graph cut methods. However, this performed poorly on some classes of brain tumors, and as a result, we modified our watershed algorithm to include a viscous component, which is comparable to simulating viscous fluid flow [8]. Our approach also involved aspects of the Watershed Cuts algorithm, a modified approach to the watershed transform in which catchment basins were generated by finding the minimum spanning trees of a graphical representation of an image [9]. The pipeline of our final approach to solving was as follows:

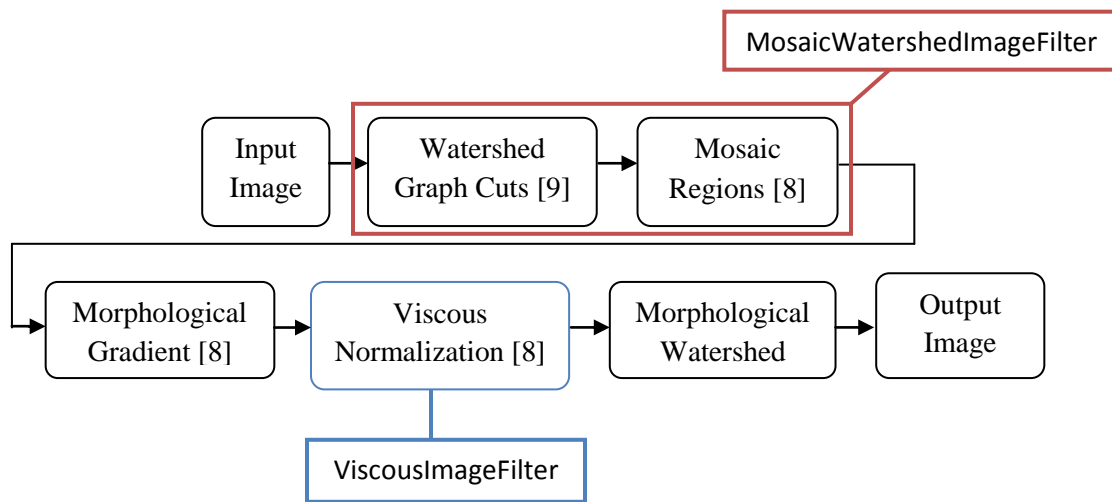


The first step, Watershed Graph Cuts, involves a watershed transform on the original input image. Rather than performing a gradient magnitude, however, we used the methods described in [9] and first found the minimum spanning trees of the image. These trees defined our regions, and were used as output to the second step in the pipeline. This step, labeled Mosaic Regions, represented the process in which each pixel in each region was set to the mean of the region. This resulted in flattened regions of homogeneous intensity. Morphological gradient was then calculated using a ‘spherical’ neighborhood, which was discretized into a 6-pixel neighborhood ± 1 pixel in the x, y, and z directions. This approximate 3-dimensional morphological gradient image was normalized by Viscous Normalization, a process in which the image is analyzed using round neighborhoods of varying radius to perform morphological reconstruction [8]. Finally, a

morphological watershed was performed on the normalized image (this method is the original watershed method in ITK). We used this method because the output of the Viscous Normalization is a gradient image, and the graph cuts method is only appropriate in the context of the original images.

This pipeline yields our output image, and differs from [8] primarily in the speed increase found in the minimum spanning tree methods of [9]. Our pipeline ultimately combines the strongest aspects of the viscous and graph cut watershed methods. Moreover, every step has been optimized for speed. During development, we tested our program for both practicality and effectiveness (see the documentation for more details).

Our pipeline (available as `ViscousWatershedImageFilter`) was designed as follows:



Implementations of all algorithms were coded in C++. C-Make and Microsoft Visual Studio 2008 were used to build ITK-SNAP, ITK, FLTK, and VTK.

Results

As mentioned above, the classical watershed algorithm involved simply a watershed transform on a gradient magnitude image. However, this yielded borders that were untrue to faint image features, and often resulted in the merging of the background regions with the borders of ring-enhancing lesions. Two typical segmentations are shown in figure 1 and figure 2. Both contain parts in the inner necrotic tissue is fused with the gray matter due to poor border definition.




Figure 1: A typical result of ring-enhancing segmentation using the classical watershed. The leak can be seen in the bottom left in the left half of the tumor.




Figure 2: A typical result of ring-enhancing segmentation using the classical watershed. The leak can be seen in the bottom left in the top right section of the tumor.




Figure 3: A typical result of ring-enhancing segmentation using the viscous watershed. We can see a loss in detail around the ring enhancement.




Figure 4: Again, we see the loss in detail around the enhancement areas.

When we replaced this with our modified pipeline, we were able to pick up the general shape of features in the image. However, the faint borders faced the issue of underrepresentation, and were often merged with both tumor and gray matter. In many cases, we were unable to separate enhancement from non-enhancement very well, as shown in figures 3 and 4. While this is problematic, we can see that with the viscous results, we are able to consistently obtain a volume in nearly all cases. In contrast, if the classical watershed leaks, then there is no way for the program to reasonably estimate volume or the general shape of the tumor.

Following this, we sought feedback from the residents. The residents each tested the program on 6 separate data sets after acclimation to the ITK-SNAP interface. While there was a noticeable speed-up in the new method, there was also an unacceptable level of accuracy loss in several scans. This, in turn, motivated our decision to pursue a post-processing step that is described in the next section.

Discussion

In our results, we showed that we were able to obtain results superior to the conventional watershed algorithm using the new hybrid viscous method. However, this came at the cost of border accuracy. A large portion of our problems stemmed from a particular subclass of tumors – the ring enhancing tumor. This tumor is characterized by a dark, necrotic center surrounded by a ring of enhancement, illustrating the selective perfusion present in the active areas of the tumor. These rings of enhancement are sometimes faint for various reasons (poor contrast administration, low scan strength, heterogeneity, etc.), and this results in vaguely defined borders that are difficult, if not impossible, for the watershed transform to recognize. As can be seen in the results section, the viscous watershed provides more round shapes for these tumors, but now fails to recognize the ring border as separate. In other words, the necrotic center is preserved (no longer merged with the gray matter) at the expense of the enhancing border. Obviously, this result will result in a shape and volume close to the actual shape and volume of the tumor. However, it will also, in nearly all cases, result in an underestimation of volume, as enhancement is most definitely representative of cancerous tissue.

We plan to address this problem with an additional post-processing step following the viscous watershed. This step will specifically address the ‘ring’ portion of the tumor. This is our current area of focus, and involves the task of separating enhancement from the area specified as cancerous by the user. Once this is done, the method is hypothesized to work on nearly all scans and classes of tumors.

In spite of this shortcoming, however, the viscous algorithm performs very robustly when compared to the regular watershed. Moreover, due to the fact that we used a fixed threshold, our program is also far less prone to inter and intra-observer variability; some other methods allow a variable threshold, which may contribute to less consistent results [10].

Conclusions

The watershed transform yields semi-consistent results in the problem of glioblastoma segmentation. However, with the viscous modification, we are able to improve the result into something that consistently yielded near-true borders. Moreover, these borders have been repeatedly shown to yield a closed border on at least the main parts of the tumor.

Future Work

Our current focus is perfection of the segmentation result via the post-processing step. This will likely be done past the semester, and into the summer.

Moreover, once the program is modified to a suitable standard, we plan to pursue a study of the software's accuracy and variability in Nijmegen, the Netherlands this summer in collaboration with Dr. Adams and the UMCN Radboud department of neurosurgery.

Acknowledgements

We would like to thank Dr. Alfredo Quiñones-Hinojosa for his support and lab space. We would also like to thank Dr. Hadie Adams for his oversight of this project, as well as Dr. Taylor for his support and valuable feedback. Finally, we would like to thank neurosurgical residents Dr. Shaan Raza, Dr. Chetan Bettagowa, and Dr. Jose Undabeitia for the feedback they provided during the development stages of the UI, and during the evaluation of the segmentation results.

References

1. Kaus, M., S. Warfield, et al. (1999). Segmentation of Meningiomas and Low Grade Gliomas in MRI. Medical Image Computing and Computer-Assisted Intervention – MICCAI'99. C. Taylor and A. Colchester, Springer Berlin / Heidelberg. **1679**: 1-10.
2. Clarke, L. P., R. P. Velthuizen, et al. (1995). "MRI segmentation: Methods and applications." Magnetic Resonance Imaging **13**(3): 343-368.
3. Keles GE, Anderson B, Berger MS: The effect of extent of resection on time to tumor progression and survival in patients with glioblastoma multiforme of the cerebral hemisphere. Surg Neurol 52:371–379, 1999.
4. Keles GE, Chang EF, Lamborn KR, Tihan T, Chang CJ, Chang SM, Berger MS: Volumetric extent of resection and residual contrast enhancement on initial surgery as predictors of outcome in adult patients with hemispheric anaplastic astrocytoma. J Neurosurg 105:34–40, 2006.
5. Lacroix M, Abi-Said D, Fourney DR, Gokaslan ZL, Shi W, DeMonte F, Lang FF, McCutcheon IE, Hassenbusch SJ, Holland E, Hess K, Michael C, Miller D, Sawaya R: A multivariate analysis of 416 patients with glioblastoma multiforme: Prognosis, extent of resection, and survival. J Neurosurg 95:190–198, 2001.
6. Pope WB, Sayre J, Perlina A, Villablanca JP, Mischel PS, Cloughesy TF: MR imaging correlates of survival in patients with high-grade gliomas. AJNR Am J Neuroradiol 26:2466–2474, 2005.
7. Yushkevich, P. A., J. Piven, et al. (2006). "User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability." NeuroImage **31**(3): 1116-1128.
8. Vachier, C. and F. Meyer (2005). "The Viscous Watershed Transform." Journal of Mathematical Imaging and Vision **22**(2): 251-267.
9. Cousty, J., G. Bertrand, et al. (2009). "Watershed Cuts: Minimum Spanning Forests and the Drop of Water Principle." Pattern Analysis and Machine Intelligence, IEEE Transactions on **31**(8): 1362-1374.
10. Letteboer, M. M. J., O. F. Olsen, et al. (2004). "Segmentation of tumors in magnetic resonance brain images using an interactive multiscale watershed algorithm1." Academic Radiology **11**(10): 1125-1138.

Management Summary

The coding of the UI was equal effort between Nathaniel and Alexander. This required building ITK, FLTK, VTK, and ITK-SNAP, and then modifying the appropriate files in all these sources for the desired result. Nathaniel worked on the parts necessary within ITK, while Alexander worked on the ITK-SNAP portions. FLTK files were done by both, and design choices were made as a group with consultation from Dr. Adams.

The minimum deliverable was to implement a two-dimensional watershed algorithm. This was met, as expected, within the expected date of March 17th. The original expected deliverables were to integrate this two-dimensional watershed algorithm into ITK-SNAP, and test both variability and accuracy of the program. While we were able to integrate the watershed algorithm into ITK-SNAP within the expected time frame, we faced difficulty with the variability study due to the inadequate 'leaky' borders encountered that were a result of the heterogeneity of the MRI scans. As a result, we modified our expected deliverables based on a literature review of our problem, and began work on a viscous variation of the watershed transform. We did this with the hopes that this modification would resolve the problem described above, while sacrificing some of the time saved.

Our new expected deliverables were changed to include the vanilla two-dimensional watershed implementation, the viscous two-dimensional watershed implementation, and finally, the algorithm documentation along with its submission to the Insight Journal. We completed the viscous C++ implementation within the expected date of May 13, along with its documentation in preparation for its submission. The results from this algorithm were not what we expected – while the new methods fixed the leaks, the problem of over-segmentation returned and there were far too many regions for the tool to be an effective time-saver.

Finally, our maximum deliverables were to evaluate the performance of this algorithm, both in terms of its accuracy and variability (both inter and intra-observer). We also planned to evaluate its usefulness and ease of use through three neurosurgical residents who had already agreed to do so. However, since our program was unable to obtain usable borders for all classes of tumors (specifically, the ring-enhancing tumors), we decided to postpone this until the summer, during which we plan to perfect the algorithm and test it in collaboration with Dr. Adams and the UMCN Radboud neurosurgery department. Thus, we changed our maximum deliverables to the documentation and publication of our algorithm, as it is the most we can expect to finish as the semester ends.

Overall, we completed both a normal and viscous watershed filter for use in ITK, along with a program that applies both filters to generate borders for segmentation of brain MRI images. We also completed the documentation for the algorithms, and are currently working on the submission. This encompasses both the minimum and expected deliverables, and most of the maximum deliverables. Next, we plan on submitting the filter to the Insight Journal and further

investigating other methods to improve the result over the coming summer. Alexander will be in the Netherlands working with Dr. Adams on the variability study, while Nathaniel will be in Baltimore working with Dr. Quiñones on the retrospective study. Towards the end of summer we plan on submitting a publication on using this program as a volumetric analysis tool.

Through this project, we have learned that the diverse spectrum of GBM tumors presents a difficult computer vision problem that most likely cannot be fully solved with simple mathematical filters. Rather, it is likely that we will finish with a volumetric tool that requires some minor manual refinement or the selection of many regions. Both will save time (the first will save much more time while sacrificing some variability, and vice versa), but the watershed transform by itself is not enough to compensate for the ambiguity in present-day MRI scans.

Note: The source code files and executables of both algorithms are available on the course website.