

Perceptual Grouping Using Active Contours in the FLUFFER Software

Geoffrey Hollinger and Dr. Bruce Maxwell
Swarthmore College
Summer 2003

Abstract

The FLUFFER (Flexible Labeler Used for Fundamental Edge Research) program developed at Swarthmore College is useful for determining image segmentations and labeling regions and edges for database construction. This paper describes an implementation of active contours with balloon forces within the software that attempts to improve the edge detection and segmentation through perceptual grouping of edges detected by the Rothwell edge detection method [1]. The fully autonomous active contour implementation is successful at grouping edge pixels in close proximity and with similar orientation into closed and partially closed contours. The active contours, however, will not seal up large holes or extrapolate edges into regions where local descriptors do not find them. The method improves FLUFFER'S segmentation of images with distinct regions and mostly linear salient features proving itself useful for image analysis.

Introduction

The FLUFFER program was developed in attempt to provide image segmentation that corresponds to coherent surfaces and gives physical characteristics of those regions as described by Maxwell [1]. To achieve this goal, the software provides tools that allow the user to label the physical characteristics of edges and regions within images. Once the labeling has been performed on many images, the output files can be examined for trends in image attributes that correspond to physical characteristics.

Corder describes the use and goals of the FLUFFER program in [3]. Her paper also describes the basic structure of the program and some preliminary research done on reflectance ratio and rate of change of intensity values using the software. FLUFFER was intended to be modular in nature and allow the easy implementation of new techniques and code. The segmentation process in the software utilizes an edge detection method, the Canny and Rothwell edge detection are currently implemented, and then a region growing segmentation method that cannot grow through the previously detected edges. The region growing can use various

colorspace to perform its tasks. The original FLUFFER software was not designed with any perceptual grouping implementations.

Boyer describes the bases of many perceptual grouping techniques in [4]. Most techniques revolve around the Gestalt theory of human psychological grouping. The Gestalt psychologists theorized that humans group objects by proximity, similarity, movement, continuity, and closure. The previous versions of FLUFFER only took advantage of similarity and proximity measurements when determining segmentation through region growing. Perceptual grouping methods such as those described in the theory sections and implemented below attempt to use continuity and closure along with proximity and similarity to provide segmentation and image analysis more similar to that of the human brain.

Zhu presents the process of grouping parts of images into a higher structure in [5]. The idea of table legs being associated as part of a table forms the basis of perceptual grouping. The possibility that a closed contour can exist within an image suggests that most edge pixels within the image can be grouped into a higher order structure or contour. Snakes, stick growing, and other perceptual grouping techniques attempt to take advantage of this relationship to close up holes in images and provide top down analysis to supplement the bottom up edge detection techniques used by Canny, Rothwell, and others. These perceptual grouping observations form the basis for my attempt to integrate active contours into FLUFFER.

Theory and Background

To perform perceptual grouping of edges, it was first necessary to use an edge detection method that finds reasonable edges with few major errors. Of the two edge detection methods implemented on FLUFFER, Canny and Rothwell, the Canny method was analyzed by Heath [6] and determined to be an excellent edge detector in his comparison test. The Rothwell method, however, produced less extraneous edges and turned out to be more useful for edge grouping. Another major advantage of using the Rothwell method was that before edges are determined, each potential edge is assigned an edge strength measurement and then thresholded. By carrying these measurements through to the perceptual grouping functions, it is possible to use them for a second thresholding procedure. The EdgeThresh active contour parameter allows the user to adjust this value. Rothwell presents the initial calculation of edge strengths in [1]. The active

contour implementation presented in this paper works with Canny edge detection, but I recommend using Rothwell edge detection to attain the full advantage.

Several popular methods of perceptual grouping have been presented to group edges into salient images features. A stick growing approach, developed by Nelson [7], groups edges by growing them together based on gradient direction, line segment orientation, and continuity. These techniques work well for the extraction of lineal features in images as shown by Nelson and Hunt [8]. While these techniques may have proved useful in grouping edges within FLUFFER, I decided to implement the more widely used active contour method described below.

The groundwork for active contours is presented by Kass in his snakes implementation. He uses an energy minimization function to guide contours into areas of regions of high gradient. The snakes maintain continuity through spring forces between "snaxels" and form a coherent contour unit. Kass's snakes are not autonomous and require user interaction to guide their initialization and analysis. To develop an autonomous active contour technique, I decided to give my initialized snakes outward forces. This idea is similar to balloons, which have been implemented in papers such as Cohen's [10]. The idea behind balloons is that an initialized contour will be given outward forces that must be offset by gradient forces to minimize the energy of the snake and stop it from continually expanding.

Implementation

To begin the implementation of a balloon like active contour, it was necessary to determine a useable data structure. My first attempt used an array of snaxel structures. The snaxel structure contained information about the location of the snaxel on the image. Since the active contour would require the quick deletion and creation of new snaxels, the array idea was scrapped due to lack of speed. The implementation currently uses a linked list of snaxels. The snake.C file contains the Contour class, which allows for the initialization, expansion, and output of an active contour of snaxels. This class contains files capable of creating, deleting, and moving snaxels within the linked list. Because of the linked list data structure, snaxel manipulation is reasonably fast.

Before contour growing, it is necessary for FLUFFER to find seeds from which to grow. I decided to allow FLUFFER to segment the images without perceptual grouping and use the

region map from this segmentation to determine contour starting points. On a faster processor, it would be possible to perform multiple iterations and use each subsequent region map for further perceptual grouping. Because of the expansion technique described below, it was sufficient to use a fairly simple method of central location inside each region of the rough segmentation as a contour seed. It was not possible, however, to use the centroid of the region because the centroid is not guaranteed to be in the interior of a region with one or more hole. Instead, the findStarts() function in contour.c determines the leftmost and rightmost pixel of each region and then scans the column between them. The function finds the bottommost pixel within the region along that column, moves upward until it leaves the region, and then uses the pixel halfway inside the region as a starting pixel. An image in which contour seeds have been found but not expanded is shown in Figure 1.

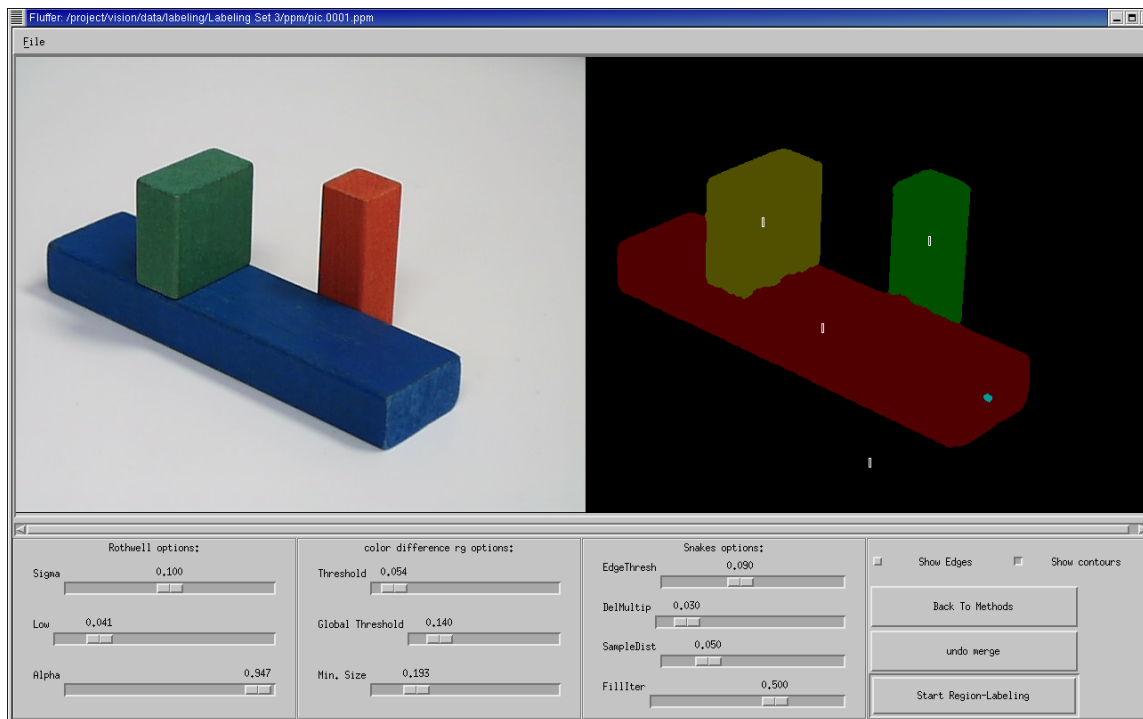


Figure 1: Snaxel seeds determined and initialized on a FLUFFER image

Once the contour seeds have been found, they need to be expanded until they have hit enough edge pixels of sufficient strength to stop the expansion. The initial implementation used a box expansion where snaxels above the starting point move up, snaxels to the left moved left and so forth. This expansion technique was only useful in finding square and rectangular contours and did not determine correct contours on irregularly shaped regions. To fix this, the expansion was then done with a circular force test around the seed coordinate. The contour

interpolated itself as it got larger interpolating pixels into each hemisphere of the circle and expanding in different regions. This technique also failed to capture irregular regions. Attempts to rotate snaxels around the seed while expanding were not feasible and too processor intensive.

The technique for contour expansion on which I finally decided is most useful for large regions with few holes. Once a contour seed has been found, a line of snaxels to the right and left of the center pixel is initialized to move away from the center. The first snaxels expand in the horizontal direction until they are forced to stop by one of the following events: they hit an edge pixel of sufficient strength to stop them, they hit the boundary of the image, or they become surrounded by enough stopped snaxels to prevent further movement. Snaxels that stop because they are surrounded by stopped pixels impose a continuity measurement within the snake. Contours cannot contain jagged edges or sharp curves because their neighbors will catch stray snaxels. This is desirable since jagged edges and curves are rare in real world images.

After the initial expansion, it is necessary to interpolate the snake and create new snaxels for vertical expansion. To do this, a line drawing function was implemented to connect snaxels. A more complicated connection algorithm such as Bezier curves or B-splines with each snaxel as a control point could have been implemented, but the line drawing was sufficient for this research. I recommend trying either of the above curves on this implementation for further research. According to a user-defined parameter, `SampleDist`, baby snaxels are created at intervals along the lines between current snaxels.

Once interpolation has been performed, it is then necessary to determine which direction these new snaxels will expand. The center of the snake is no longer useful because interpolation has created snaxels around any holes in the region. To determine the interior of the contour, it is necessary to fill the boundaries created by connecting the snaxels. To do this, the snaxels are connected by edge lines and the resulting boundary is filled using a seed fill algorithm [11]. A scan line or flood fill algorithm would probably be quicker but also more complicated. This is recommended for further research. Once the interior of the contour has been filled, it can now be tested to determine expansion direction for baby snaxels. After direction is determined, each snaxel is moved two pixels in that direction to prevent missing edges continuous with surrounding snaxels. This process continues where babies switch off moving horizontally and then vertically on different iterations. This repeats until the snake has expanded as many times as defined by `FillIter`, another parameter.

After the snaxel expansion and interpolation had finished, I noticed that the contours were connecting many stray pixels to the areas of edges that were correctly connected. Fixing this required a two-step process. First, as part of the `cleanUpSnake()` function, all snaxels that were not sitting on edge pixels were deleted. This deleted any snaxels that were stopped due to continuity infringement or ones that had been newly interpolated. Once extraneous snaxels had been deleted, it was then necessary to correct improper connections. The simplest way to do this was to disconnect snaxels that were beyond a certain multiple of `SampleDist` apart. The `DeleteMult` parameter gives the multiple of `SampleDist` at which disconnection occurs. This cleanup prevents most contours from being closed but, as can be seen in Figures 2 and 3, greatly decreases extraneous connections. I also tried disconnecting snaxels that created angular discontinuities within the snake, but this did not provide as pleasing of a cleanup.

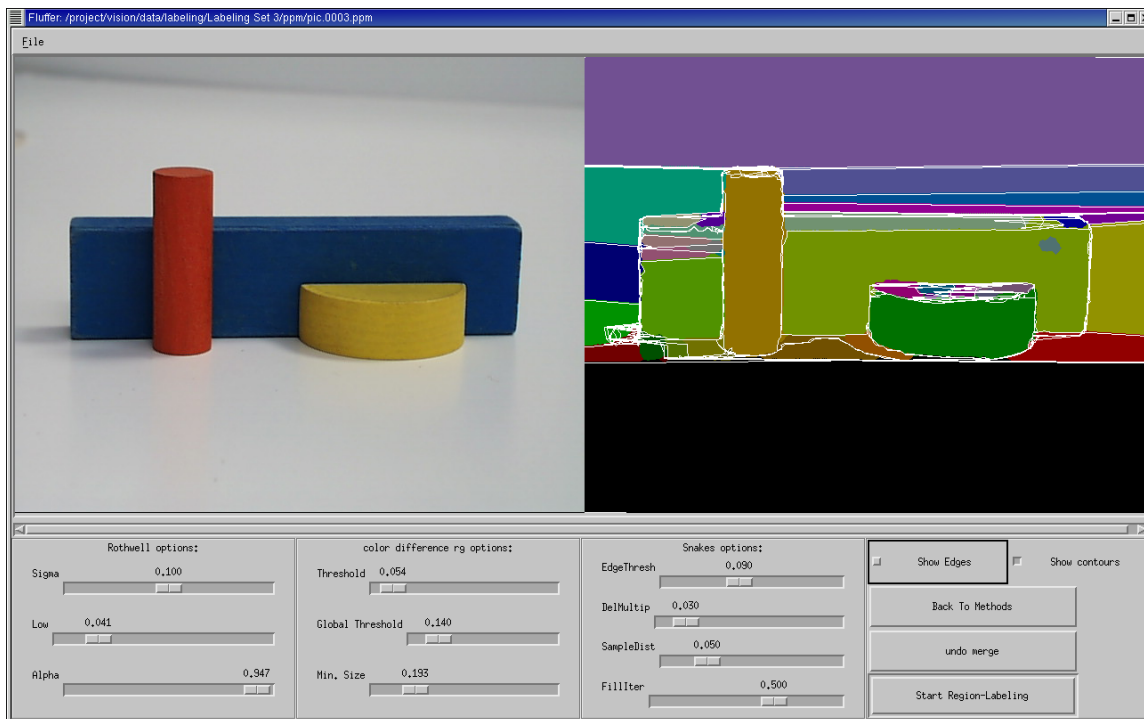


Figure 2: Contour determination before cleanup function

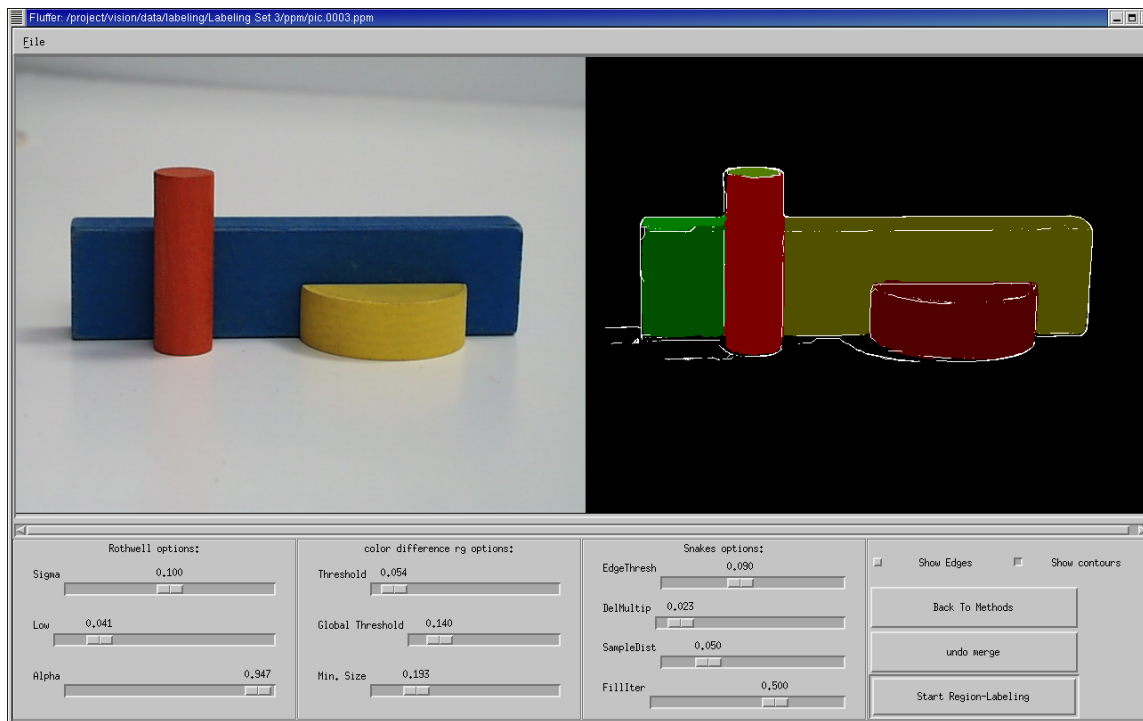


Figure 3: Contour determination after cleanup function

Advantages and Limitations

The implementation of active contours described in this paper has certain advantages and disadvantages for various images and region types. It does an excellent job of grouping together edges in close proximity on the outside of distinct regions. The perceptual grouping often seals edges that are not correctly grouped together by the Rothwell edge detector. Active contours are also useful for finding intra-object edges that are not found by the region growing technique. If the local edges are distinctly defined, the perceptual grouping will often seal them up enough that the region growing will not bleed through. This creates distinct regions between intra-object edges. Such grouping is especially desirable in FLUFFER because of its intended use for determining trends among various edge types.

Some disadvantages of this active contours implementation include its failure to close up jagged edges. Because of otherwise useful continuity restraints, the snakes will not expand into tight corners of amorphous regions. Also, incorrect expansion can lead to connections across non-edge boundaries that do not correspond to actual edges. This is a danger in any non-local edge detection or perceptual grouping technique and is mostly compensated for by the cleanup function in this case. Finally, the active contours in this paper cannot group edges across large

holes that lack local edges. The human brain can use shape and object recognition to seal up these edges, while the active contours do not have that advantage.

Conclusions and Further Research

In conclusion, autonomous active contours provide a useful supplement to FLUFFER's segmentation process. Their proficiency in sealing up intra-object edges allows the classification of edges that would not be found using the region growing and one of the bottom up edge detection processes. They do, however, have some notable disadvantages and double FLUFFER's operating time for most images. Further research to improve perceptual grouping by active contours is recommended below.

Prominent research in image segmentation, such as that done by Liao [12], has used B-splines to represent active contour connections and continuity forces. Such a change in the internal structure of the snakes would likely result in tighter fitting to edge boundaries and more correct connections. It also would provide a new method of error detection by examining the stresses among the spline. Furthermore, researchers such as N. Xu [13] have experimented with growing contours to a width greater than a single pixel and then shrinking them into high gradient areas. This may prove useful as a further cleanup exercise after the deletion and disconnection has been done. The remaining contours could be settled into high gradient areas to better gather data on edge characteristics.

The current active contour implementation does not use any raw image characteristics. Characteristics of any of the colorspace used by FLUFFER could be used to influence snaxel expansion and contour formation. One particularly useful technique is described by C. Xu [14] and uses Gradient Vector Flow (GVF) forces to guide contour deformation. GVF forces are promising in that they allow snaxels to be guided into small crevices in regions and to be attracted to areas of high gradient from large distances. It may be possible to guide snaxels from the start using GVF forces changing the implementation considerably. Calculation of GVF forces, however, is mathematically complex and processor intensive and will likely cause a considerable operating time hit.

Sources

[1] Rothwell, C., J. Mundy, W. Hoffman and V.D. Nguyen. "Driving Vision by

- Topology." Proceedings IEEE Symposium on Computer Vision, pp. 395-400, 1995.
- [2] Maxwell, Bruce and Steven A. Shafer. "Segmentation and Interpretation of Multicolored Objects with Highlights." Computer Vision and Image Understanding, Vol. 77, pp. 1-24, 2000.
- [3] Corder, Julie. "FLUFFER: Edge Analysis for Computer Vision Research." Swarthmore College, Swarthmore, PA, Spring 2003.
- [4] Boyer, Kim L. and Sudeep Sarkar, ed. Perceptual Organization for Artificial Vision Systems. Boston: Kluwer Academic, 2000.
- [5] Zhu, Song-Chun. "Statistical Modeling and Conceptualization of Visual Patterns." IEEE Trans. on PAMI, Vol. 25, No. 6, June 2003.
- [6] Heath, Mike, et al. "Comparison of Edge Detectors." Computer vision and Image Understanding, Vol. 69, No. 1, pp. 38-54, 1998.
- [7] Nelson, Randal C. "Finding Line Segments by Stick Growing." IEEE Trans. on PAMI, Vol. 16, No. 5, May 1994.
- [8] Hunt, Galen C. and Randal C. Nelson. "Lineal Feature Extraction by Parallel Stick Growing." University of Rochester, Rochester, NY.
- [9] Kass, M., A. Witkin and D. Terzopoulos. "Snakes - Active Countour Models." International Journal of Computer Vision, 1(4): 321-331, 1987.
- [10] Cohen, L.D. and I. Cohen. "Finite-Element Methods for Active Contour Models and Balloons for 2-D and 3-D Images." IEEE PAMI, Vol. 15, No. 11, pp. 1131-1147, 1993.
- [11] Rogers, David F. Procedural Elements for Computer Graphics. 2nd ed. Boston: McGraw-Hill, 1979.
- [12] Liao, Chia-Wei and Gerard Medioni. "Simultaneous Surface Approximation and Segmentation of Complex Objects." Computer Vision and Image Understanding, Vol. 73, No. 1, January, pp. 43-63, 1999.
- [13] Xu, Ning, Ravi Bansal and Narendra Ahuja. "Object Segmentation Using Graph Cuts Based Active Contours." IEEE CVPR Proc., Vol. 2, pp. 46-53, 2003.
- [14] Xu, C. and J.L. Prince. "Gradient Vector Flow: A New External Force for Snakes." Proc. IEEE Conf. on Comp. Vis. Patt. Recog. (CVPR). LosAlamitos: Comp. Soc. Press, pp. 66-71, June 1997.
- [15] Stokman, H.M.G. and Th. Gevers. "Classification of Color Edges in Video into Shadow-

Geometry, Highlight, or Material Transitions." University of Amsterdam, Amsterdam, The Netherlands.

Appendix A: Other Changes Made to FLUFFER

A.1: Simple Parameterization of Edge Detection and Segmentation

As a further improvement to FLUFFER, I modified the graphical user interface to allow for a simpler way of choosing edge detection and segmentation parameters. When simple parameter mode is activated for either or both processes, the user is given a single scroll bar labeled sensitivity for that process. When the user changes the sensitivity, FLUFFER then determines the original three parameters for segmentation and edge detection as a linear function of the sensitivity value. A high sensitivity will result in many edges and many regions, while a low sensitivity will result in few edges and few regions.

The same linear mapping was used to give the user the option of simple perceptual grouping parameters when utilizing active contours. These modifications move towards a parameter free edge detection and segmentation in which parameters are determined autonomously at the local level. This has not yet been achieved in FLUFFER but is presented by Stockman in source [15]. Because of the simplicity to the user and potential speedups in edge labeling, I recommend further research in autonomous parameter determination for FLUFFER. Using a more complex mapping of the sensitivity parameter onto the original parameter space could also improve simple parameterization. At this point, however, because of the complexity of the parameters for all methods, I urge the use of complex parameterization when labeling a database for analysis.

A.2: Option to Ignore Small Regions

On a dual processor Pentium II 400 MHz system, FLUFFER takes approximately ten seconds to segment a 512x640 ppm image without perceptual grouping. The most time consuming process in the software is the small region handling algorithm that reassigns regions that are smaller than a user defined size and assigns them to similar larger regions. The perceptual grouping implementation described in this paper requires the region growing and small region handling process to be done twice. To avoid greatly increasing processing time, I

implemented a faster small region handling technique for the segmentation done before perceptual grouping.

To improve small region handling speed for the rough run through, I had FLUFFER assign all small regions to the same region. When it was determining the starting points for the contour seeds, it ignored any region that was assigned to the small region value. This allowed FLUFFER to perform the same segmentation mentioned above on the same processor in less than 20 seconds with perceptual grouping. To make this change, it was necessary to separate the `Image::handleSmallRegions()` function from `Image::completeSegmentationGR()` and add a new function `Image::ignoreSmallRegions()` to the `Image` class. These changes were done in `twopass.C`. I also gave the user the option of using the ignore small region technique for the final segmentation. If this is used, the edge labeler ignores any edges between small regions and defined regions. While this is useful, I recommend using the original small region handling for the final segmentation to achieve a more complete edge labeling of the final image.

A.3: Minor Fixes and Code Changes

In order to expand the modular capabilities of FLUFFER, I split the original segmentation options into colorspace and segmentation categories. The colorspace options include all of the options present in the previous version. By splitting segmentation into a separate category, however, it was possible to give the user the option of using the two-pass segmentation algorithm, which was already present in `twopass.C`. This algorithm, however, proved to be almost useless when compared to the region growing algorithm, and I do not recommend it for research.

To further improve the modularity of FLUFFER, I moved the code for reading the configuration files from `methodPanel.c` to a new file called `methodConfig.c`. This file reads in all configuration data for the edge detection and segmentation methods as well as newly necessary data for the colorspace, small region handling, and perceptual grouping methods. This file also contains the code to extract the simple parameterization. Any changes to this I/O should be done by modifying the new `methodConfig.c` file. Within this file, a minor bug was fixed that assigned memory for segmentation attributes based on the maximum number of edge attributes. If a segmentation method were implemented with more than three attributes, the previous version of FLUFFER would have crashed.

After image segmentation has been performed in FLUFFER, it is sometimes necessary to manually merge regions by first clicking on one region and then clicking on the other. While doing this, it was sometimes confusing as to which region the user had clicked on. To help fix this, I added a function `highRegionInSegmentDa()` to the `imageDrawing.c` file. This function is called when a region is clicked on for merging. It sets all pixels in that region to white so that the user can determine exactly which region he or she has selected.