Evolutionary Optimization of Filter Parameters for Image Segmentation

Timothy Rupe[1] Jennifer Leopold[1] Anne Maglia[2] Daniel Tauritz[1]

[1]Department of Computer Science, University of Missouri-Rolla
[2]Department of Biological Sciences, University of Missouri-Rolla

Abstract

This paper explores the use of evolutionary algorithms (EA) for parameter selection of image segmentation algorithms. Typically, segmentation algorithms are tuned "by hand" by the user through modifying various combinations of parameters – a time consuming and computationally expensive process. EA provide a means to explore possible parameter combinations without user trial-and-error. Herein we show the advantages of applying EA to segmentation algorithms using an example application of 3D medical image reconstruction. We also highlight future improvements to our method that will increase its efficiency and usability.

Keywords: Evolutionary algorithms, image segmentation, 3D visualizations

1. Introduction

This paper explores the use of evolutionary algorithms (ES) to select parameters for existing image segmentation algorithms. The algorithms partition an image into disjoint regions that follow two basic properties: 1) the regions must be homogeneous, encompassing data that share common characteristics such as color and texture, and 2) they must be heterogeneous relative to one other, having distinct boundaries. Formally, the problem is described in the following way [2]. Let I be an image and H be a homogeneity predicate; then the segmentation of I is a partition of I into a set of regions $\{R_1, R_2, ..., R_n\}$ such that:

$$\cup_{i=1}^{n} R = I,\ R_i \cap R_j = \varnothing,\ i \neq j$$

1.1. Background

Technological developments in the field of medical imaging are increasing at an exponential rate, and hundreds of researchers and clinicians are generating three-dimensional (3D) visualizations of anatomical structures from two-dimensional (2D) MRI, CT and histological sections. 3D visualizations have revolutionized the study of anatomy, and provide researchers and clinicians with new tools to diagnose disease, train surgeons and examine internal structures.

Most current anatomical visualization technologies involve rendering a single 3D volumetric reconstruction that can be sliced along the x, y, or z axis (e.g., Figure 1a), but cannot be partitioned into 3D sub-objects. A more biologically useful strategy is to segment the 3D reconstruction into sub-objects representing the various anatomical parts that comprise the real biological object (Figure 1b).

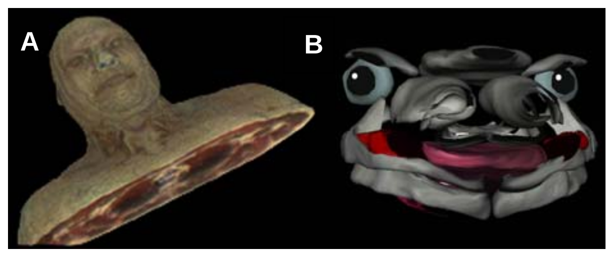

Fig. 1: A) Sliced reconstruction of human (Laboratory of Integrated Systems, University of São Paulo, Brazil); B) Segmented reconstruction of frog head (*MorphologyNet* digital library, University of Missouri-Rolla).

Unfortunately, current image segmentation software is inefficient and requires a great deal of user input to increase the accuracy of the results. Two factors limit the optimization of image segmentation algorithms. First, the number of possible combinations of parameters and filters is astronomical. Thus, finding an optimal (or near optimal) solution by evaluating every possible permutation or performing random searches is unacceptable. Second, each image is unique, and an acceptable set of parameters for one image may not be acceptable for another.

1.2. Goal

Our goal is to apply evolutionary algorithms (EA) to image partitioning so as to minimize the number of parameter-tuning iterations required to find a set of parameters that result in an acceptable image (based on user opinion). Evolutionary algorithms allow for a directed search of potential parameter sets and allow for incremental steps toward better combinations, while requiring little human guidance (in the form of grading the algorithm's performance).

2. EA Design and Implementation

Most tasks involving image processing require source images to have similar properties. Typically, if input images vary, custom filters must be created for different images. A better strategy would be to create a general purpose solution with stochastic operations.

Evolutionary algorithms are ideal for creating a general purpose solution because they allow for dynamic searching of optimal filter properties [3]. In addition, they use a fitness function that provides direction to the search (thus avoiding inefficient random searches). Human intervention is a common means of providing guidance and training to an EA [4]. Because the optimization of most image processing tasks can be highly subjective, human intervention is particularly necessary for this problem. However, great care must be taken to limit the amount of human input, so as to improve the overall efficiency of the process.

The EA is not directly used to perform image processing, but rather is used to tune the parameters of the various filtering algorithms [5]. This allows traditional image processing algorithms to perform the raw manipulation, while the EA guides the overall process. In the experiments described herein, we simplified the overall process by using the same image filters and algorithms for all trials, and applying the EA only to parameter selection. This limits the overall complexity of the experiments, while allowing us to examine the effects of using the EA. Future experiments will include the application of EA to filter selection.

Fortunately, the Visible Human Project (VHP), hosted by the National Library of Medicine, provides many powerful image segmentation tools, such as the Insight Visualization Toolkit (ITK), that are ideal for testing the current problem. The ITK provides filters and constructs for segmenting both 2D and 3D images, as well as a plug-in interface for new custom filters. For our experiments, we developed a method to dynamically modify the ITK filters to optimize segmentation quality through the application of EA.

2.1. Input Parameter Files

Our parameter files contain a list of unique problem descriptions, where each line describes the desired

P	O	T	Input	Output
10	20	5	thorax.bmp	thorax_result.bmp
15	30	15	neck.bmp	neck_result.bmp

Table 1. Example problem descriptions. P = population size, O = number of offspring, T = tournament size.

parameters for a unique run (e.g., Table 1). Each line contains the type of information found in Table 1.

2.2. Input Data Files

Each input data file is a standard bitmap image (Figure 2) in one of several common formats (e.g., jpeg, bmp, tiff). For our experiments, we used bmp files because they are uncompressed and simple in structure.

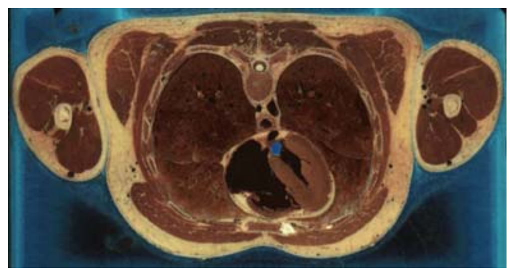

Fig. 2: Example input file: bmp of human thorax (cross section) from the Visible Human Server.

2.3. Output Data Files

Output data is represented as a bitmap image file, with distinct colored shapes depicting individual regions identified by the segmentation algorithms (Figure 3). The output format defaults to that of the input image, or can be specified by file extension.

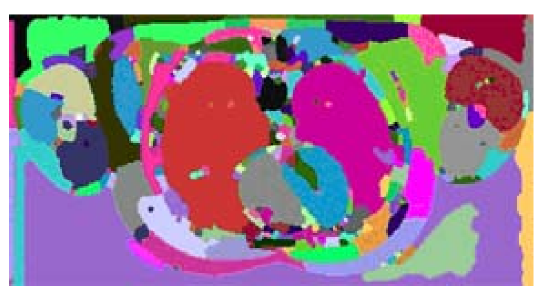

Fig. 3: Example output file: bmp of segments identified in Figure 2

2.4. Representation

Each individual in the segmentation problem represents a set of parameters for the image processing filters. A common representation of an individual is described as an array of real values, the first half of which represent filter parameters and the second half represents the standard deviation parameters used in EA uncorrelated mutation with n step sizes, where n represents the number of parameters. Our representation differs slightly to accommodate the input parameters expected by the image manipulation filters – i.e., floating point, integer, and boolean values. The use of boolean values negates the need to use a mutation parameter because a boolean value is never "partially" changed. Examples of two randomly generated individuals are shown in Table 2.

C	Its	T	Scale	G	s_1	s_2	s_3	s_4
float	integer	float	float	boolean	Float			
2.0	10	0.0	0.05	1	1.0	1.0	1.0	1.0
2.0	10	0.001	0.15	0	1.0	1.0	1.0	1.0

Table 2. Example of individual representation. C = conductance, T = threshold, G = gradient.

2.5. Fitness Function

The segmentation quality of an image is highly subjective – i.e., the same solution can score very well or very poorly depending on user opinion. We evaluated the fitness of each candidate based on points and age, which allow more flexible image evaluation, but require human intervention to award points.

During the evaluation phase (described below), the user gives points to individual solutions believed to be better than others presented. As the EA progresses through each generation, individuals that survive will accumulate point scores and age. The fitness of an individual is defined by its score/age ratio:

$$Fitness = \frac{score}{age}$$

The goal is to maintain only individuals with a high fitness ratio. Individuals who age, but do not score well, will have a lower fitness, while individuals who receive high points from generation to generation will increase in age, but will maintain a high fitness value. This method allows high scoring individuals a high probability of survival to the next generation, while also allowing new individuals to compete.

2.6. Population Initialization

To begin the trials, a small initial population of candidate solutions (about 10) is created and the resulting segment images are presented to the user. The user chooses the candidate solutions to keep, which are then used to create children that are added to the population. Using this guided initialization, the first population has a higher average fitness than a random sample, and thus, limits the number of poor performers that the user must evaluate.

2.7. Parent Selection

For each child to be created, two parents are selected from the population through tournament selection with tournament size k (i.e., for a given tournament, k individuals are randomly chosen from the population, and the best is selected as a parent). The size of the tournament can be adjusted to increase or decrease the selective pressure. For λ children, 2λ parents must be selected.

2.8. Recombination

Child creation from two parents is performed in two stages. First, the parameter values for the child are filled by local recombination. For each parameter in the child, the value is chosen randomly from one or the other parents until all parameter values are filled. Second, the standard deviation values of the child are filled using global recombination – i.e., using the population average for that value.

2.9. Mutation

Mutation of children is based on uncorrelated mutation with n step sizes (as borrowed from the concept of evolutionary strategies), and is performed in two steps: 1) standard deviation values σ are mutated, and 2) based on the new values, the parameter values x are mutated. Formally, the standard deviation values and parameter values are mutated in the following way [1]:

$$\sigma' = \sigma_i \cdot e^{\tau' \cdot N(0,1) + \tau \cdot N_i(0,1)}$$

$$x'_i = x_i + \sigma_i \cdot N_i(0,1)$$

Note that because one of the parameters is a boolean value, there is no σ value associated with it; thus, it is mutated by randomly flipping it with a probability of $1/n$.

2.10. Evaluation

The parameters of an individual and input image are passed to the image segmentation algorithm. The results are visualized as the graphical representation of the identified segment regions. The resulting image is evaluated by the user, and if the user feels positively toward the result, the individual representing the image is rewarded points. It is important to note that the same individual may not be awarded points from one generation to the next because another image may perform better or the user may simply change his/her mind.

2.11. Survivor Selection

To minimize the amount of human interaction, it is important to use high selective pressure. Survivor selection uses a deterministic rank-based method, utilizing the (μ,λ) method, wherein all children are included in the population, and the lowest ranking individuals are automatically removed.

2.12. Termination Condition

Because there is no target fitness value, the optimal output depends entirely on the user. The evolutionary cycle continues indefinitely until the user finds a satisfactory result or resigns after finding only unsatisfactory results.

3. Experiments and Results

Several different trials with increasing image complexity using the EA implementation discussed above were conducted. All trials used the same starting parameters: population = 10, offspring = 20, and tournament size = 5.

3.1. Simple Shapes

This trial tested the algorithm's ability to segment an image consisting of three regions (Figure 4 "input"). As expected, the algorithm found all three regions easily (Fig. 4 "output").

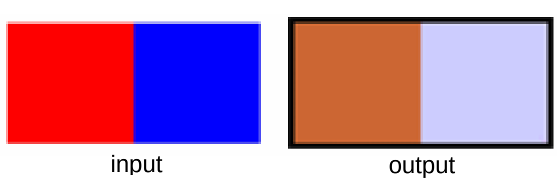

Fig. 4. Results of simple image experment. All regions were identified. (Note that input object has a white boarder).

3.2. Moderately Complex Shapes

This trial used more complex images with multiple regions of different shapes and sizes (Fig. 5 "input"). Results were quite good, but the results show very small discrepancies in recognizing the tips of the star (Fig. 5 "output), likely because of the sharp angles in those areas.

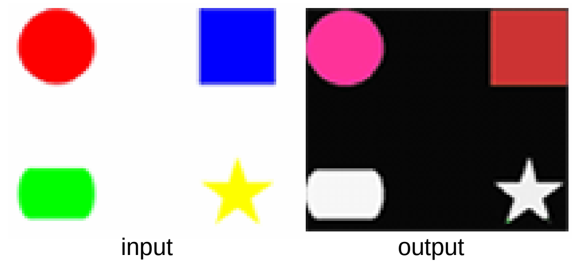

Fig. 5. Results of a moderately complex shape experiment. All regions were identified, but slight errors recognizing the tips of the stars were evident.

3.3. Complex Shapes

This trial included more complex and overlapping shapes (Fig 6. "input"). The only difficulty for the algorithm was in identifying small regions such as outlined borders (Fig. 6 "output"). Instead of a line being represented as a single shape, it is broken into small individual regions. However, the resulting segmentation quality was still good, and this trial was terminated after three generations.

3.4. Biological Data

This trial used the image in Figure 2 as input. This image is very complex, with varying colors, textures, and potential region sizes. After thirteen generations, the user felt the results were acceptable (Fig. 3), and the trial was terminated. The segmentation was not perfect, but most of the primary body tissues were identified, and only a small amount of noise was present.

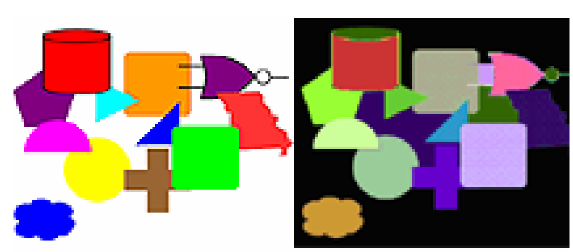

Fig. 6. Results of complex shape experment. All regions were identified, but slight errors recognizing the borders of the objects were evident.

4. Future Work

There is a number of design modifications that we plan to make to the implementation. First, in designing the experiments, we simplified the overall process by using the same image filters and algorithms for all trials, and only applied an EA to parameter selection. In future experiments, we will also apply an EA to filter selection. Secondly, when evaluating the fitness of an individual, the user currently must award a fixed number of points, amounting to a pass/fail system. In future implementations, the user will have more flexibility in awarding points. Thirdly, current output is in graphical format, but in future implementations output will also be saved as vector data so that it can be imported into 3D modeling programs.

In addition, we plan to develop a graphical user interface for presenting image data to the user. Currently the user must open each image file to view the data. We hope to integrate the user-evaluation process with the tournament selection process and present the user with a group of images from which to choose.

Finally, we hope to accelerate the process by using faster segmentation algorithms and by reducing the amount of user-intervention through the use of target templates. If a user could provide a template image of the ideal result, the program could be modified to automatically compare the fitness function of an individual to the target's fitness function, thereby reducing the number of alternatives the user must evaluate.

5. Conclusions

The current implementation required about one hour of processing time per generation. In addition, each individual in a generation had to be evaluated by the user. Thus, the amount of time and user involvement required to process the data prevent the current implementation from being used as is.

However, overall the results of the trials were excellent, indicating that the application of an EA improves the efficiency and accuracy of image segmentation algorithms. Thus, with the above future modifications to the current implementation, we are confident that we will develop a powerful tool for automating image processing.

6. References

[1] A. Eiben and J. Smith, *Introduction to Evolutionary Computing*. Natural Computing Series. Springer-Verlag, Berlin Heidelberg, 2003.

[2] C. Lai and C. Chang, "A Hierarchical Genetic Algorithm Based Approach for Image Segmentation", *Networking, Sensing and Control, 2004 IEEE International Conference, Vol.2*, pp. 1284-1288, 2004.

[3] C. Munteanu, V. Lazarescu, and C. Radoi, "A New Strategy in Optimization Using Genetic Algorithms", *Proc. of Electrotechnical Conference, 1998. MELECON 98, 9th Mediterranean*, Volume 1, pp. 415-419, May 1998.

[4] C. Munteanu and A. Rosa, "Evolutionary Image Enhancement with User Behavior Modeling", *SIGAPP Applied Computing Review*, 9(1):8-14, 2001.

[5] W. Smart and M. Zhang, "Applying Online Gradient Descent Search to Genetic Programming for Object Recognition", Proceedings of the Second Workshop on Australasian Information Security, Data Mining and Web Intelligence, and Software Internationalization, pp. 133-138. Australian Computer Society, Inc., 2004.