From: AAAI Technical Report FS-93-04. Compilation copyright © 1993, AAAI (www.aaai.org). All rights reserved.

Adaptive Image Segmentation Using Multi-Objective Evaluation and Hybrid Search Methods

Bir Bhanu, Sungkee Lee, and *Subhodev Das*
College of Engineering
University of California
Riverside, CA 92521-0425

Abstract

This paper describes an approach for image segmentation that relies on learning from experience to adapt and improve the segmentation performance. The adaptive image segmentation system incorporates a feedback loop consisting of a machine learning subsystem, an image segmentation algorithm, and an evaluation component which determines segmentation quality. The machine learning component is based on genetic adaptation and uses (separately) a pure genetic algorithm (GA) and a hybrid of GA and hill climbing (HC). When the learning subsystem is based on pure genetics, the corresponding evaluation component is based on a vector of evaluation criteria. For the hybrid case, the system employs a scalar evaluation measure which is a weighted combination of the different criteria. Experimental results for pure genetic and hybrid search methods are presented using a representative database of outdoor TV imagery.

1 INTRODUCTION

Image segmentation is an important and, perhaps, the most difficult low-level task. The difficulty arises when the segmentation performance needs to be adapted to the changes in image quality. Image quality is affected by variations in environmental conditions, imaging devices, time of day, etc. Despite the large number of segmentation techniques presently available [4], no general methods have been found that perform adequately across a diverse set of imagery. When presented with a new image, selecting the appropriate set of algorithm parameters is the key to effectively segmenting the image [2]. However, no segmentation algorithm can automatically generate an "ideal" segmentation result in one pass (or in an open loop manner) over a range of scenarios encountered in real-world applications. Any technique, no matter how "sophisticated" it may be, will eventually yield poor performance if it can not adapt to the variations in unstructured scenes.

In reality, there exist several factors which make the parameter adaptation process very difficult. *First*, the number of parameters present in a typical segmentation algorithm is usually quite large. *Second*, the parameters mutually interact in a complex, non-linear fashion, which makes it difficult or impossible to model their behavior in an algorithmic or rule-based fashion. *Third*, since variations between images cause changes in the segmentation results, the objective function that represents segmentation quality also varies from image to image. *Finally*, the definition of the objective function itself can be a subject of debate because there is no single, universally accepted measure of segmentation performance available with which to uniquely define the quality of the segmented image.

Consequently, there exists a need to develop an adaptive segmentation technique that can efficiently search the complex space of plausible parameter combinations and locate the values which yield optimal results. The approach should not be dependent on the particular application domain nor should it have to rely on detailed knowledge pertinent to the selected segmentation algorithm. While there are adaptive threshold selection techniques for segmentation, these techniques do not accomplish any learning from experience to improve the performance of the system over time. In the absence of any rigorous theory, the problem of image segmentation is best described in terms of its goal. The criteria for good segmentation are [4], (1) the segmented regions should be uniform and homogeneous with respect to some characteristic, such as gray value or texture, (2) region interiors should be free of holes and region boundaries should be smooth and spatially accurate, and (3) adjacent regions should be differing significantly based on the characteristic on which they are uniform. If one represents this criteria set in terms of a function, then the problem of (good) segmentation is one of optimizing this objective function by selecting appropriate segmentation parameters.

2 ADAPTIVE SEGMENTATION ALGORITHM

Adaptive image segmentation requires the ability to modify control parameters in order to respond to changes that occur in the image as a result of varying environ-

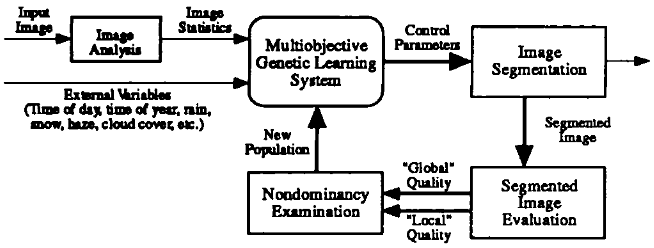

Figure 1: Block diagram of the adaptive image segmentation system for multiobjective optimization.

mental conditions. The block diagram of our approach to adaptive image segmentation is shown in Figure 1. After acquiring an input image, the system analyzes the image characteristics and passes this information, in conjunction with the observed external variables, to the machine learning component (GA or GA-HC hybrid). Using this data, the machine learning system selects an appropriate parameter combination, which is passed to the image segmentation process. After the image has been segmented, the results are evaluated and an appropriate reward is generated and passed back to the learning subsystem. This process continues until a segmentation result of acceptable quality is produced.

The image segmentation component in our work is the Phoenix algorithm [5, 6] which has been extensively tested on color imagery. Phoenix contains seventeen different control parameters [5] fourteen of which are used to control the thresholds and termination conditions of the algorithm. There are 10^{33} conceivable parameter combinations using these fourteen values. Of the fourteen values, we have selected two of the most critical parameters that affect the overall results of the segmentation process: *maxmin* and *hsmooth*. From an analysis of the Phoenix algorithm, we find that incorrect values in the two main parameters lead to results in which, at one extreme, the desired object is not extracted from the background, and at the other extreme, the object is broken up into many small regions that have little significance for higher-level processes. By measuring segmentation performance using appropriate quality criteria, the genetic process attempts to identify a parameter set that yields results between these two extremes.

2.1 Multi-Objective Segmentation Evaluation

In order to overcome the drawbacks of using only a single quality measure [7], we have incorporated an evaluation technique that uses five different quality measures described below to determine the overall fitness for a particular parameter set. These are,

1. *Edge-Border Coincidence:* Measures the overlap of the region borders in the segmented image with the edges found in the original image using an edge operator.
2. *Boundary Consistency*: Similar to edge-border coincidence, except that region borders which do not exactly overlap edges can be matched with each other. In addition, region borders which do not match with any edges are used to penalize the segmentation quality.
3. *Pixel Classification*: This measure is based on the number of object pixels classified as background pixels and the number of background pixels classified as object pixels.
4. *Object Overlap*: Measures the area of intersection between the object region in the groundtruth image and the segmented image.
5. *Object Contrast*: Measures the contrast between the object and the background in the segmented image relative to the object contrast in the groundtruth image.

The maximum and minimum values for each of the five segmentation quality measures are 1.0 and 0.0, respectively. The first two quality measures, i.e., edge-border coincidence and boundary consistency, are ***global*** measures since they evaluate the segmentation quality of the whole image with respect to edge information. Conversely, the last three quality measures are ***local*** measures since they only evaluate the segmentation quality for the object regions of interest in the image. When an object is broken up into smaller parts during the segmentation process, only the largest region which overlaps the actual object in the image is used in computing the local quality measures.

The three local measures require the availability of groundtruth information in order to correctly evaluate the segmentation quality. Since groundtruth data may not always be available, the adaptive segmentation system is designed to use three separate methods of evaluating segmentation quality. First, segmentation quality can be measured using global evaluation method alone. Second, if groundtruth data is available and we are only interested in correctly segmenting the object regions in the image, then the local evaluation method can be used alone. Finally, if we desire good object regions as well as high quality overall segmentation results, then the global and local quality measures together can be used to obtain a ***scalar-valued*** or a ***vector-valued*** segmentation quality measure that maximizes overall performance of the system. The maximization of the vector-valued segmentation quality measure is in effect a multiobjective optimization problem where the global and the local measures represent the "non-commensurable" criterion functions.

A multiple objective constrained optimization problem is of the form

$$\max[f_i(\mathbf{x}) = z_i], i = 1, ..., k, \text{such that } \mathbf{x} \in S \quad (1)$$

where $f_i(\mathbf{x})$'s are the objective functions and z_i's are the corresponding optimal criterion values and S is the feasible region. However, it is only in the trivial case, that there exists a single point in S which simultaneously maximizes all k objectives. A typical approach in multiobjective (or vector-valued) optimization is to consider the *utility* of the z_i's. Thus, a point in S is optimal if it maximizes the decision maker's utility function. To be optimal, however, a point must be *efficient* or *Pareto* optimal.

The key concept of Pareto optimality is the "partially greater than" ($p>$) relation between two vectors of the same dimension. Given two vectors $\mathbf{a} = (a_1, ..., a_n)$ and $\mathbf{b} = (b_1, ..., b_n)$, $\mathbf{a}$ is said to be partially greater than $\mathbf{b}$ ($\mathbf{a}\, p> \mathbf{b}$) if each element of $\mathbf{a}$ is greater than or equal to the corresponding element of $\mathbf{b}$ and at least one element of $\mathbf{a}$ is strictly greater than the corresponding element of $\mathbf{b}$, i.e.,

$$(\mathbf{a}\, p> \mathbf{b}) \Longleftrightarrow (\forall i)(a_i \geq b_i) \;\; (\exists i)(a_i > b_i).$$

Under these conditions, we say that $\mathbf{a}$ dominates $\mathbf{b}$ or $\mathbf{b}$ is inferior to $\mathbf{a}$. If a vector is not dominated by any other vector, it is said to be *nondominated* or non-inferior.

A description of the multiobjective optimization approach for image segmentation using a pure GA is given below.

1. Compute the image statistics.
2. Generate an initial population.
3. Segment the image using initial parameters.
4. Compute the global and local quality measures.
5. Examine nondominancy of each individual.
6. WHILE not <stopping conditions> DO
   - 6a. select subgroups of individuals using each dimension of the quality measures
   - 6b. generate new population using the crossover and mutation operators
   - 6c. segment the image using new parameters
   - 6d. compute global and local quality measures
   - 6e. examine nondominancy of each individual

   END
7. Update knowledge base using new structures.

The stopping criteria for the multiobjective optimization system consist of two conditions. First, the process terminates if an utopian parameter set, i.e., the one for which both local and global quality measures are above a predefined threshold of acceptance, is located. The thresholds for acceptable segmentation is 90% of the best segmentation. This criterion is useful only when the best for each segmentation quality surface is known *a priori*. Second, the process terminates if both the average local quality and the average global quality of the populations decrease for three consecutive generations or fail to improve for five consecutive generations. If either of these conditions is met, the segmentation of the current image is stopped and the nondominated parameter sets are represented as the current best estimates of the Pareto-optimal set.

2.2 Hybrid Search Combining Genetic Algorithm and Hill Climbing

Genetic algorithms have been proven and shown to provide robust search performance across a broad spectrum of problems [3]. However, hybrid techniques [1] have the potential for improved performance over single optimization techniques since these can exploit the strengths of the individual approaches in a cooperative manner. One such hybrid scheme which is the focus of this paper combines a global search technique (genetic algorithm) with a specialized local search technique (hill climbing). Hill climbing (HC) methods are not suitable for optimization of multimodal objective functions, such as the segmentation quality surfaces, since they only lead to local extrema and their applicability depends on the contour shape of the objective functions. The hybrid scheme provides performance improvements over the genetic algorithm alone by taking advantage of both the genetic algorithm's global search ability and the hill climbing's local convergence ability. In a sense, the genetic algorithm first finds the hills and the hill climber climbs them.

The block diagram of the adaptive image segmentation system using the hybrid optimization scheme is shown in Figure 2. and the description is given below.

1. Compute the image statistics.
2. Generate an initial population.
3. Segment the image using initial parameters.
4. Compute the segmentation quality measures.
5. WHILE not <stopping conditions> DO

   IF <new maximum found>
   - 5HCa. generate all points (i.e., parameters) adjacent to the current point
   - 5HCb. segment image using these points
   - 5HCc. compute the quality measures
   - 5HCd. climb to new maximum point if it exists

   ELSE
   - 5GAa. select individuals using reproduction
   - 5GAb. generate new population using the crossover and mutation operators
   - 5GAc. segment the image using new parameters
   - 5GAd. compute segmentation quality measures

   END
6. Update knowledge base using new structures.

3 EXPERIMENTAL RESULTS

A database of 20 outdoor images of a static scene, collected over a 5-hour period, is used by the system. Figure 3(a) shows one selected frames of the database. The car in the image is the object of interest for the pixel classification, object overlap, and object contrast

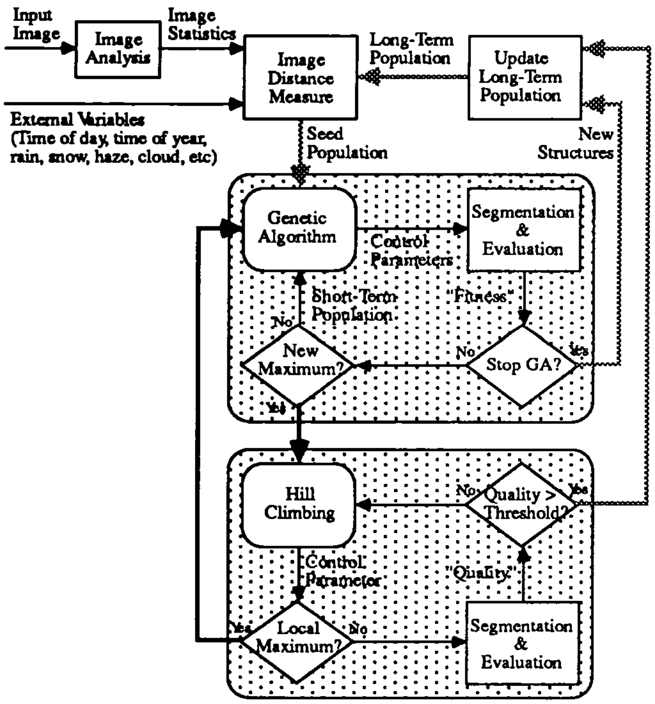

Figure 2: Block diagram of the adaptive image segmentation system using the hybrid search scheme.

segmentation quality measures. The ground truth image for the car is obtained by manual segmentation of Frame 1 (not shown) of the image sequence. The segmentation quality surfaces, both global and local, for each frame is exhaustively defined for preselected ranges of maxmin and hsmooth parameters of the Phoenix algorithm. Default values are used for the remaining parameters. Figures 3(b)-(c) show the global and local quality surfaces for the frame of Figure 3(a). The ten odd-numbered images are selected as the training data, while the remaining even-numbered images are used for testing. The genetic component uses a long-term population size of 100 individuals, a short-term population size of 10, a crossover rate of 0.8, and a mutation rate of 0.01. The stopping criteria is 90% of the global and local maxima of the global and local quality surfaces of each image in the database.

Each training image is processed 100 times, each with a different (randomly selected) seed population. The search points visited on the quality surface at the various generations while processing a training image (Frame 3) during multi-objective optimization are shown in Figure 4(a), e.g., points at the lower-hand corner of the graph correspond to global and local segmentation quality of 0.0. Figure 4(b) displays the utopian point at the upper right corner, which caused the termination of the genetic search process after third generation. In this figure, segmentation performance over 90% is denoted as 100%. Figure 4(c) displays the segmentation result for Frame 3. This result was obtained from the individual in the short-term population with maximum local fitness (i.e., the best local segmentation quality for the car). During testing, the seed population is selected from the long-term population obtained at the end of the training experiments. Since the fitness values of the testing seed population are usually high, the GA converged to the Pareto-optimal set much faster during the testing experiments than in the training experiments.

The same outdoor imagery database is used for the hybrid algorithm. Also, the training and testing sequences are kept unchanged. Recall that the fitness is now a scalar, combining the global and local segmentation quality measures, for each individual of the genetic population. To provide a visual indication of the performance improvements achieved by the adaptive segmentation system using the hybrid search scheme, the segmented image results are shown for Frame 3 in Figure 5. These results are obtained using the individual from the short-term population that has the maximum fitness (i.e., correspond to the best segmentation quality). Each of these segmented images shows a tendency to obtain more precise boundary representations for each of the background objects as well as the border of the car. The hybrid scheme results surpassed the pure GA-based results in 8 (out of 10) training images when the number of segmentations required to optimize the segmentation quality was reduced. On the average, an improvement of 15.3% in performance was observed with the training images. However, the hybrid scheme-based testing results are no better than the pure genetics results, because the training results supplied the testing seed points located in highly fit regions of the search space which could hardly be optimized by the hill climbing process.

4 CONCLUSIONS

The multiobjective optimization demonstrates the ability of the adaptive image segmentation system to provide high quality segmentation results in a minimal number of generations. The results of the hybrid method show the performance improvement over the GA alone. In general, the hybrid scheme performs better than the pure GA for the frames which require less computational effort to optimize the segmentation quality, i.e., for the frames which have simpler segmentation quality surfaces.

References

[1] D. Ackley. Stochastic iterated genetic hillclimbing. Technical Report CMU-CS-87-107, Dept. Computer Science, Carnegie Mellon University, Pittsburgh, Pa., March 1987.

[2] B. Bhanu, S. Lee, and J. Ming. Self-optimizing image segmentation system using a genetic algorithm. In *Plenary Presentation, Fourth Intl. Conf. Genetic Algorithms*, pages 362–369, San Diego, CA, July 1991.

(a)

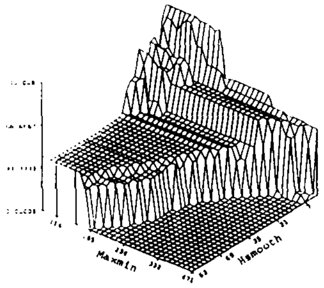

(b)

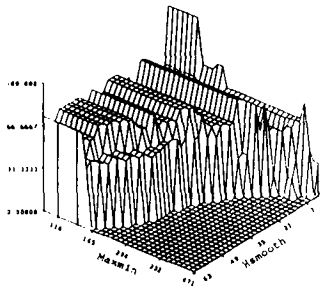

(c)

Figure 3: A representative color image for the adaptive segmentation algorithm: (a) Frame 3. Segmentation quality surfaces: (b) Global, (c) Local for Frame 3.

[3] D. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Reading, Mass.: Addison-Wesley, 1989.

[4] R. M. Haralick and L. G. Shapiro. Image segmentation techniques. *Comp. Vis. Graph. Image Proc.*, 29:100–132, 1985.

[5] K. I. Laws. The phoenix image segmentation system: Description and evaluation. Technical Report 289, SRI International, Dec. 1982.

[6] R. Ohlander, K. Price, and D. R. Reddy. Picture segmentation using a recursive region splitting method. *Comput. Graphics and Image Processing,* 8:313–333, 1978.

[7] J. Schaffer and J. Grefenstette. Multi-objective learning via genetic algorithms. In *Proc. 9th Intl. Joint Conf. Artificial Intelligence*, pages 593–595, Los Angeles, Calif. 1985.

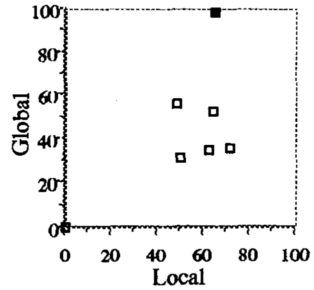

(a) Generation 1

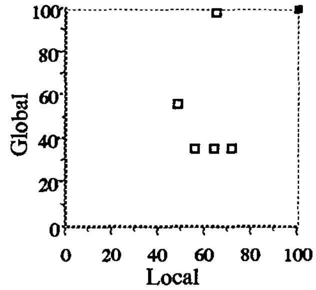

(b) Generation 3

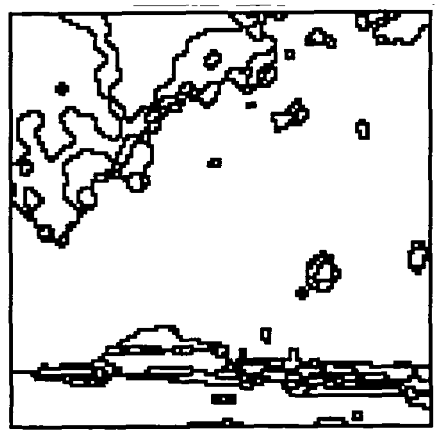

(c)

Figure 4: Segmentation (training) results for Frame 3: (a) global and (b) local segmentation quality of each individual at each generation, and (c) local segmented image. The dark squares represent the locally nondominated points at each generation.

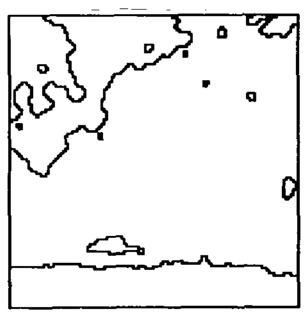

(a) Generation 1

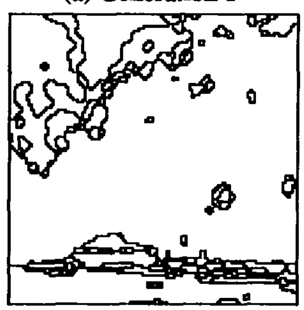

(b) Generation 6

Figure 5: Segmented images corresponding to Frame 3 in the hybrid experiment.