Self-Optimizing Image Segmentation System Using a Genetic Algorithm

Bir Bhanu
College of Engineering
University of California
Riverside, CA 92521

Sungkee Lee
Dept. of Computer Sc.
University of Utah
Salt Lake City, UT 84112

John Ming
Human Interface Tech. Ctr.
NCR Corp., 500 Tech Parkway
Atlanta, GA 30313

Abstract

One of the fundamental weaknesses of current computer vision systems to be used in practical outdoor applications is their inability to adapt the segmentation process as real-world changes occur in the image. We present the first closed loop image segmentation system which incorporates a genetic algorithm to adapt the segmentation process to changes in image characteristics caused by variable environmental conditions such as time of day, time of year, clouds, etc. The segmentation problem is formulated as an optimization problem and the genetic algorithm efficiently searches the hyperspace of segmentation parameter combinations to determine the parameter set which maximizes the segmentation quality criteria. The goals of our adaptive image segmentation system are to provide continuous adaptation to normal environmental variations, to exhibit machine learning capabilities, and to provide robust performance when interacting with a dynamic environment. We present experimental results which demonstrate that genetic algorithm can be successfully used to adapt the segmentation performance automatically in outdoor color imagery.

1. INTRODUCTION

Image segmentation is typically the first, and most difficult task of any automated image understanding process. All subsequent interpretation tasks including object detection, feature extraction, object recognition, and classification rely heavily on the quality of the segmentation process. Despite the large number of segmentation techniques presently available [3,6], no general methods have been found that perform adequately across a diverse set of imagery. Only after many modifications to its control parameter set can any current segmentation technique be used to process the diversity of images found in real world applications. When presented with a new image, selecting the appropriate set of algorithm parameters is the key to effectively segmenting the image. The image segmentation problem can be characterized by several factors which make the parameter selection process very difficult. *First*, most segmentation techniques contain numerous control parameters which must be adjusted to obtain optimal performance. The size of the parameter search space in these systems can be prohibitively large, unless it is traversed in a highly efficient manner. *Second*, the parameters within most segmentation algorithms typically interact in a complex, non-linear fashion, which makes it difficult or impossible to model the parameters' 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 varies from image to image. The search technique used to optimize the objective function must be able to adapt to these variations. *Finally*, the definition of the objective function itself can be subject to debate because there are no universally accepted measures of image segmentation quality.

Hence, we must apply a technique that can efficiently search the complex space of 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. The key elements of our adaptive image segmentation system are: (1) A closed-loop feedback control technique that consists of a genetic learning component, an image segmentation component, and a segmented image evaluation component; (2) A genetic learning system that optimizes segmentation performance of each image and accumulates segmentation experience over time to reduce the effort needed to optimize the segmentation quality of succeeding images; (3) Image characteristics and external image variables are represented and manipulated using both numeric and symbolic forms within the genetic knowledge structure, only the segmentation parameters are represented and manipulated in binary strings; (4) Image segmentation performance is evaluated using five measures of segmentation quality that measure *global* characteristics of the entire image as well as *local* features of individual object regions; (5) The adaptive segmentation system is not dependent on any specific segmentation algorithm or type of sensor. The performance of the adaptive algorithm will be limited by the capabilities of the segmentation algorithm, but the results will be optimal for a given image based on our evaluation criteria.

To date, no segmentation algorithm has been developed which can automatically generate an "ideal" segmentation result in one pass (or in an open loop manner) over a range of scenarios encountered in practical outdoor applications. While there are adaptive threshold selection techniques

[12,14] for segmentation, these techniques do not accomplish any learning from experience to improve the performance of the system. Any technique, no matter how "sophisticated" it may be, will eventually yield poor performance if it can not adapt to the variations in outdoor scenes. Therefore, in this paper we attempt to address this fundamental bottleneck in developing "useful" computer vision systems for practical scenarios by developing a closed-loop system that incorporates a genetic algorithm and automatically adapts the segmentation algorithm's performance by changing its control parameters and will be valid across a wide diversity of image characteristics and application scenarios.

2. ADAPTIVE IMAGE SEGMENTATION SYSTEM

2.1 SEGMENTATION AS A SEARCH PROBLEM

Fig. 1 shows an outdoor image and the typical segmentation quality surface (discussed in Section 2.2.4) associated with the image in which only two segmentation parameters [8,13] are being varied. Because of the large number of potential parameter combinations and the subtle interaction of the algorithm parameters, the objective function is complex, multimodal, and presents problems for many commonly used search and optimization techniques. The drawbacks to some of these methodologies for the segmentation optimization problem are summarized by Lee [9].

Genetic algorithms [1,2,4,5,7] which are designed to efficiently locate an approximate global maximum in a search space show great promise in solving the parameter selection problem encountered in the image segmentation task. Since they use simple recombinations of existing high quality individuals and a method of measuring current performance, they do not require complex surface descriptions, domain specific knowledge, or measures of goal distance. Moreover, due to the generality of the genetic process, they are independent of the segmentation technique used, requiring only a measure of performance (which we refer to as segmentation quality) for any given parameter combination.

2.2 ADAPTIVE IMAGE SEGMENTATION

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 environmental conditions. The block diagram of our approach is shown in Fig. 2. After acquiring an input image, the system analyzes the image characteristics and passes this information, in conjunction with the observed external variables, to the genetic learning component. Using this data, the genetic 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 genetic algorithm. This process continues until a segmentation result of acceptable quality is produced. The details of each component in this procedure will be described in the following subsections.

2.2.1 Image Characteristics

We compute twelve first order image properties for each color component (red, green, and blue) of the image. These features include mean, variance, skewness, kurtosis, energy, entropy, x intensity centroid, y intensity centroid, maximum peak height, maximum peak location, interval set score, and interval set size. Since we use a black/white version of the image to compute edge information and object contrast during the evaluation process, we also compute the twelve features for the Y (luminance component) image as well. Combining the image characteristic data from these four components yields a list of 48 elements. In addition, we utilize two external variables, time of day and weather conditions. The external variables are represented symbolically in the list structure (e.g., time = 9am, 10am, etc. and weather conditions = sunny, cloudy, hazy, etc). The distances between these values are computed symbolically when measuring image similarity. The two external variables are added to the list to create an image characteristic list of 50 elements. A system consisting of knowledge structures is used to store the image characteristics and the associated segmentation parameters that are generated by the genetic learning system.

2.2.2 Genetic Learning System

Fig. 3 shows a simple example of our genetic learning system. The image characteristics for a new image are compared with the individuals in the global population to obtain the initial seed for the local population. The global population represents the accumulated segmentation experience for all images that the system has processed whereas the local population contains the set of segmentation parameters processed by the genetic algorithm during the optimization of the current image. To obtain the initial local population (seed population) for a new image from the global population, a normalized Euclidean feature distance is computed from the new image to every member of the global population and this distance is used along with the fitness of each individual in the global population for selecting the closest individuals. Although we have limited the seed population to 3 in this example, our experiments utilize a seed population of 10 individuals. The global population holds 100 knowledge structures in order to maintain a diverse collection of segmentation experience. The parameter sets in the seed population are used to segment the image and the results are evaluated to generate a fitness for each individual. The fitness value (leftmost value in the list) varies from 0.0 to 1.0 and measures the quality of the segmentation parameter set. Note that only the fitness value and the action portion (segmentation parameters) of the knowledge structure are subject to genetic adaptation; the conditions (image characteristics) remain fixed for the life of the structure. If the fitness values are not acceptable, the individuals are recombined and the process repeats. Each pass through the loop (segmentation-evaluation-recombination), is known as a generation. The cycle continues until the maximum fitness achieved at the end of a generation exceeds some threshold. The global population is updated using the high quality members of the local population from the current image and the system is then ready to process another image.

2.2.3 Segmentation Algorithm

Since we are working with color imagery in our experiments, we have selected the well known Phoenix segmentation algorithm developed at Carnegie Mellon University [8,10,13]. Phoenix, which was the subject of several PhD dissertations, has been widely used, refined, and documented. The algorithm, which is based on a recursive region-splitting approach, has been extensively tested on color imagery. Phoenix [8] contains fourteen different control parameters which are used to control the thresholds and termination conditions used within 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*. Maxmin specifies the lowest acceptable peak-to-valley-height ratio used when deciding whether or not to split a large region into two or more smaller parts. Hsmooth controls the width of the window used to smooth the histogram of each image region during segmentation. The use of only two parameters for the initial tests aids in the ***visualization*** of the optimization process since we can plot the associated segmentation quality corresponding to each parameter combination using a 3D plotting technique. Future research will incorporate a larger number of modifiable parameters.

2.2.4 Segmentation Evaluation

There are a large number of segmentation quality measures that have been suggested, although none have achieved widespread acceptance as a universal measure of segmentation quality. In order to overcome the drawbacks of using only a single quality measure, we have incorporated an evaluation technique that uses five different quality measures to determine the overall fitness for a particular parameter set. The five segmentation quality measures are:

(1) *Edge-Border Coincidence:* Measures the overlap of the region borders in the image acquired from the segmentation algorithm relative to an edge image.
(2) *Boundary Consistency:* Similar to edge-border coincidence, except that region borders which do not exactly overlap edges can be matched with each other. Also, region borders which do not match with any edges are used to penalize the segmentation quality [9].
(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 ground truth 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 ground truth 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 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 object ground truth information in order to correctly evaluate segmentation quality. Since we desire good object regions as well as high quality overall segmentation results, we have combined global and local quality measures (with equal weighting) to obtain a *combined* segmentation quality measure that maximizes overall performance of the system. Fig. 4 shows the surfaces defined for the five individual quality measures that are used to create the combined quality measure surface shown in Fig 1.

3. EXPERIMENTAL RESULTS

An initial database of outdoor imagery was collected to demonstrate the system's ability to adapt to real world conditions and produce the best segmentation result based on our evaluation criteria. The database consists of twenty frames that were collected approximately every 15 minutes over a 5 hour period (1:30 pm to 6:30 pm) using a JVC GXF700U color video camera. A representative subset of these images is shown in Fig. 5. This database will be used to describe the experimental results. Weather conditions in our image database varied from bright sun to overcast skies. Varying light level is the most prominent change throughout the image sequence, although the environmental conditions also created varying object highlights, moving shadows, and many subtle contrast changes between the objects in the image. The car in the image is the object of interest. The auto-iris mechanism in the camera was functioning, which causes a similar appearance in the background foliage throughout the image sequence. Even with the auto-iris capability built into the camera, there is still a wide variation in image characteristics across the image sequence. This variation requires the use of an adaptive segmentation approach to compensate for these changes.

To precisely evaluate the effectiveness of the adaptive image segmentation system, we exhaustively defined the segmentation quality surfaces for each frame. The segmentation quality surfaces were defined for preselected ranges of maxmin and hsmooth parameters Maxmin values, which affect segmentation performance in a non-linear fashion, were sampled exponentially over a range of values from 100 to 471. Values near 100 were spaced closer together than values at the upper end of the range. Hsmooth values were sampled linearly using numbers between 1 and 63. By selecting 32 discrete values (5 bits of resolution) for each of these parameter ranges, the search space contained 1024 different parameter combinations.

3.1 BASIC EXPERIMENTS

The first set of experiments with the adaptive segmentation system was divided into two separate phases: 1) a training phase where the optimization capabilities of the genetic algorithm were measured; and 2) a testing phase where we evaluated the reduction in effort achieved by utilizing previous segmentation experience. The image data was separated into two halves, 10 images (1,3,...,19) for training and 10 images (2,4,...,20) for testing. During the training phase, seed populations were selected using random locations on the combined segmentation quality surface for each image. The genetic system was then invoked using the seed population for each image and the convergence rate of the process was measured. Each training image was processed 100 times, each with a different collection of random starting points. These results were combined to compute the average number of generations needed to optimize each surface. The genetic component used a local population size of 10, a crossover rate of 0.8, and mutation rate of 0.01. A crossover rate of 0.8 indicates that, on average, 8 out of 10 members of the population will be selected for recombination during each generation. The mutation rate of 0.01 implies that on average, 1 out of 100 bits is mutated during the crossover operation to insure diversity in the local population.

The stopping criteria for the genetic process contains three tests. First, since the global maximum for each segmentation quality surface was known a priori (recall that the entire surface was precomputed), the first stopping criteria was the location of a parameter combination with 95% segmentation quality or higher. In experiments where the entire surface is not precomputed, this stopping criteria would be discarded. Second, the process terminates if 3 consecutive generations produce a decrease in the average population fitness for the local population. Third, if 5 consecutive generations fail to produce a new maximum value for the average population fitness, the genetic process terminates. If any one of these three conditions is met, the processing of the current image is stopped and the maximum segmentation quality currently in the local population is reported.

Fig. 6 shows the combined segmentation quality surfaces for the images shown in Fig. 5. Note that due to the complexity of these surfaces, most commonly-used search techniques [9] would not be effective at optimizing the segmentation quality.

At the end of training phase, the final local population from each of the training images (1,3,...,19) was combined to create a global population of 100 individuals. From this global population, the 10 initial seed members of each local population for the testing images (2,4,...,20) were selected. The testing was performed in a parallel fashion; the final local population for each of the testing images was not placed back into the global population for these tests. The alternative approach to testing, which processes each frame in the outdoor imagery database in a sequential manner and integrates the results into the global population, has been discussed in detail in [9].

For testing phase since the fitness of each seed population is based on previous segmentation experience, the genetic process is able to converge to the global maximum much faster during the testing phase. During the training experiments, the maximum number of generations was 13, the minimum number was 5, and the average number of generations was 9. By combining the information accumulated during training in the global population, the average number of generations was reduced from 9 during training to 3 during testing. This represents a considerable improvement in the adaptive system's efficiency. On average, the adaptive segmentation system visits approximately 2.5% of the search space (i.e., ~ 2.5 generations) for the experiments described here.

Since there are no other known adaptive segmentation techniques in the computer vision field to compare our system with, we measured the performance of the adaptive image segmentation system relative to the set of default Phoenix segmentation parameters [8,13] and a traditional optimization approach. The *default parameters* have been suggested after extensive amounts of testing by various researchers who developed the Phoenix algorithm [8]. The parameters for *traditional approach* are obtained by manually optimizing the segmentation algorithm on the first image in the database and then utilizing that parameter set for the remainder of the experiments. This approach to segmentation quality optimization is currently standard practice in state-of-the-art computer vision systems. Fig. 7 presents the comparison of these three approaches. The average segmentation quality for the adaptive segmentation technique was 95.8% (average of 100 experiments). In contrast, the performance of the default parameters was only 55.6% while the traditional approach provided 63.2% accuracy. As the figure shows, the performance of both of these alternative approaches was highly erratic throughout the sequence. Fig. 8 illustrates the quality of the segmentation results associated with the adaptive system, the default parameters, and the traditional approach. Each result corresponds to the average segmentation performance produced by each technique for the first frame in the database. By comparing the extracted car region in each of these images, as well as the overall segmentation of the entire image, it is clear that the adaptive segmentation results are superior to the other methods.

3.2 COMPARISON OF THE ADPATIVE SYSTEM WITH RANDOM SEARCH

Several tests were performed to compare the optimization capabilities of the adaptive segmentation system with a simple random walk through the search space. This experiment used only the training images (1,3,...,19) from the outdoor image database so that the adaptive system would not benefit from the reuse of segmentation experience from one image to the next. The intent of this restriction was to measure the efficiency of the genetic algorithm in optimizing a complex surface. In addition, the stopping criteria for the adaptive system was simplified so that when a surface point with 95% segmentation quality or better was located, the optimization process would terminate. The random walk algorithm searched the segmentation quality surface by visiting points randomly and used the same 95% stopping criteria. Finally, in order to insure correctness of the results, each segmentation quality surface was optimized by each technique 100 times and the results are averaged to create the performance figures.

Fig. 9 presents a comparison of the efficiency for the two techniques described above. The bars represent the total number of points visited on the surface using each technique for each of the images and the average number of points visited for each approach. As the average values show, the adaptive technique is far superior to the random walk approach. In addition, the average number of points visited by the adaptive approach is ~ 6.9% of the total number of points on the surface, compared to the earlier experiments where we processed ~ 2.5% of the surface, since we have not reused any segmentation experience gained from processing earlier images. Fig. 10 contrasts the segmentation quality achieved by the two techniques. Since the adaptive segmentation technique insures the achievement of a near global maximum for each image, we modified the random walk approach so that it would terminate after the same number of visited locations required by the adaptive technique. The maximum segmentation quality achieved by the random approach was then compared with the adaptive system. On the average, the adaptive system achieved 99.3% segmentation quality after the number of segmentations shown in Fig. 9. In comparison, the random walk achieved only 81.4% of the maximum quality for the same number of segmentations for each image.

3.3 THE EFFECTIVENESS OF THE REPRODUCTION AND CROSSOVER OPERATORS

A number of tests were performed to demonstrate the effectiveness of the reproduction and crossover operators in the adaptive image segmentation system. The optimization capability of the pure genetic algorithm was compared with two variations of the genetic algorithm. The first variation of the pure genetic algorithm was implemented without a reproduction operator. Instead of reproducing individuals according to their fitness values, the algorithm selected the individuals at random for further genetic operator action with the restriction that any individual be selected only once. The second variation of the genetic algorithm simply skipped a crossover operator. To ensure that this approach generates about the same number of offsprings as the pure genetic algorithm, the mutation rate of this approach was increased to the crossover rate (0.8) of the genetic process. The stopping criteria for each technique is to locate a surface point with 95% or higher segmentation quality. In order to ensure correctness of the results each image was tested by each technique 100 times and the results were averaged to create the performance figures. Fig. 11 presents the comparison of the optimization capability for three techniques. As the histograms show, the pure genetic algorithm results are much better than the results of the other two approaches for both the training and testing experiments. This demonstrates that the reproduction and crossover operators are critical for the success of genetic algorithms.

4. CONCLUSIONS

We have shown the ability of the adaptive image segmentation system using genetic algorithm to provide high quality (> 95%) segmentation results in a minimal number of segmentation cycles. The performance improvement provided by the adaptive system was consistently greater than ~33% over the traditional approach or the default segmentation parameters [8,13]. Using outdoor data, for the first time, we have shown that learning from experience can be used to improve the performance of the segmentation process. There are many more experiments that we have performed when the segmentation quality is a vector valued function [11] and optimization technique is either a pure genetic algorithm or a combination of genetic algorithm and hill climbing [9]. Although the segmentation and interpretation processes are interlinked, we have investigated how much the segmentation performance can be improved without complicating the adaptive segmentation process with the effects of recognition system performance. There are many ways in which adaptive image segmentation system described here can be used in practical computer vision systems. These research topics are currently under investigation.

Acknowledgements

The authors are grateful to Honeywell Inc. where part of this work was performed under an initiative grant.

References

1. K. DeJong, "Learning with Genetic Algorithms: An Overview," *Machine Learning* 3 pp. 121-138 (1988).

2. J.M. Fitzpatrick, J.J. Grefenstette, and D. Van Gucht, "Image Registration by Genetic Search," Proceeding of IEEE Southeastern Conference, pp. 460-464 (1984).

3. K.S. Fu and J.K. Mui, "A Survey on Image Segmentation," *Pattern Recognition* **13** pp. 3-16 (1981).

4. A.M. Gillies "Machine Learning Procedures for Generating Image Domain Feature Detectors," Ph.D. Thesis, dept. of Computer and Communication Sciences, University of Michigan, (April 1985).

5 D.E. Goldberg, "Computer-Aided gas Pipeline Operation Using Genetic Algorithms and Rule Learning," Ph.D. Thesis, Dept. of Civil Engineering, University of Michigan (1983).

6. R.M. Haralick and L.G. Shapiro, "Image Segmentation Techniques," *Computer Vision, Graphics and Image Processing* **29** pp. 100-132 (1985).

7. J.H. Holland, "Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems," pp. 593-623 in *Machine Learning: An Artificial Intelligence Approach, Vol. II,* ed. R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, Morgan Kaufmann Publishers, Inc.)1986).

8. K.I. Laws, "The Phoenix Image Segmentation System: Description and Evaluation," SRI International Technical Note No. 289 (December 1982).

9. S.K. Lee, "Adaptive Image Segmentation," Ph.D. Thesis, Dept. of Computer Science, University of Utah (July 1990).

10. R. Ohlander, K. Price, and D.R. Reddy, "Picture Segmentation Using a Recursive Region Splitting Method," *Computer Graphics and Image Processing* **8** pp. 313-333 (1978).

11. J.D. Schaffer, "Multiple Objective Optimization and Vector Evaluated Genetic Algorithms," Proceedings of International Conference on Genetic Algorithms and Their Applications, pp. 93-100 (1985).

12. P.G. Selfridge, "Reasoning About Success and Failure in Aerial Image Understanding," Ph.D. Thesis, University of Rochester (1982).

13. S. Shafer and T. Kanade, "Recursive Region Segmentation by Analysis of Histograms," Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 1166-1171 (1982).

14. J.S. Weszka, "A survey of Threshold Selection Techniques," *Computer Graphics and Image Processing* **7** pp. 259-265 (1978).

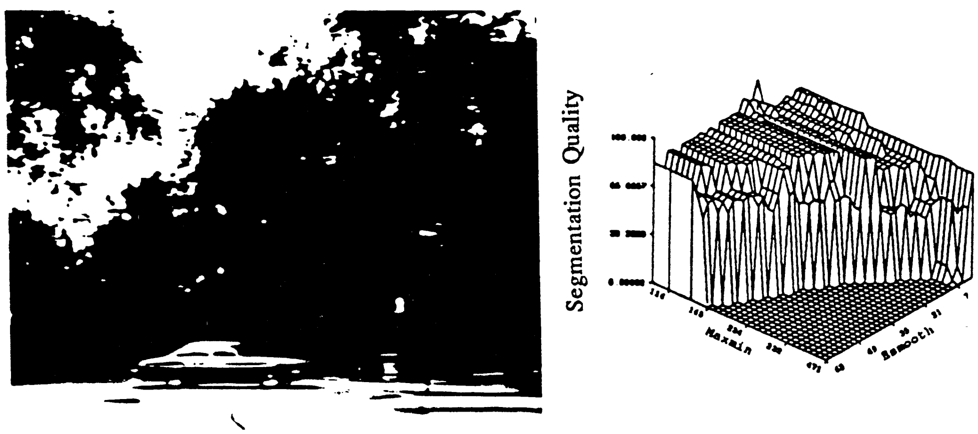

(a) (b)

Figure 1: An outdoor image and its associated segmentation quality surface. *(a)* Frame 1 of outdoor image database. *(b)* Combined segmentation quality surface.

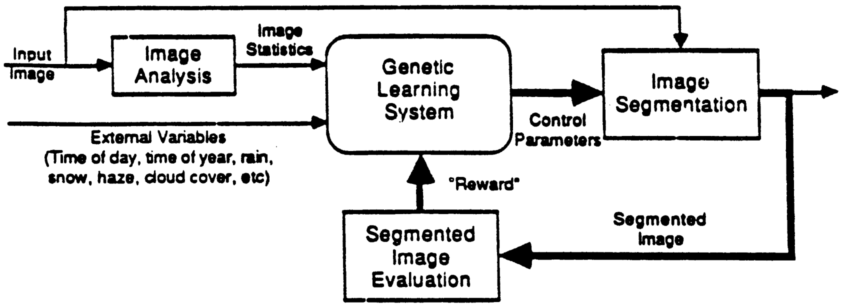

Figure 2: Block diagram of the adaptive image segmentation system.

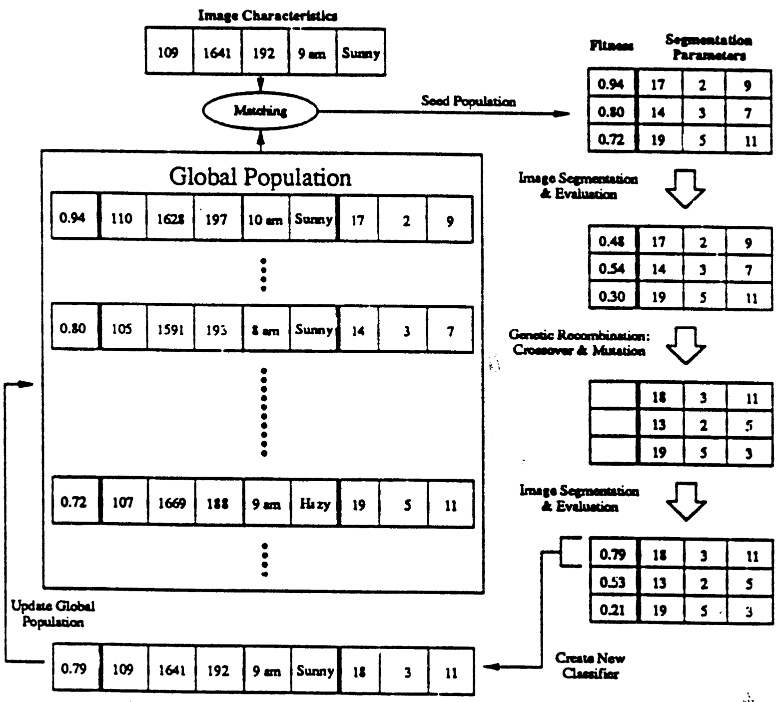

Figure 3: One complete cycle through the adaptive image segmentation

Frame 13

Frame 20

Figure 5: Selected color images from the outdoor experiments.

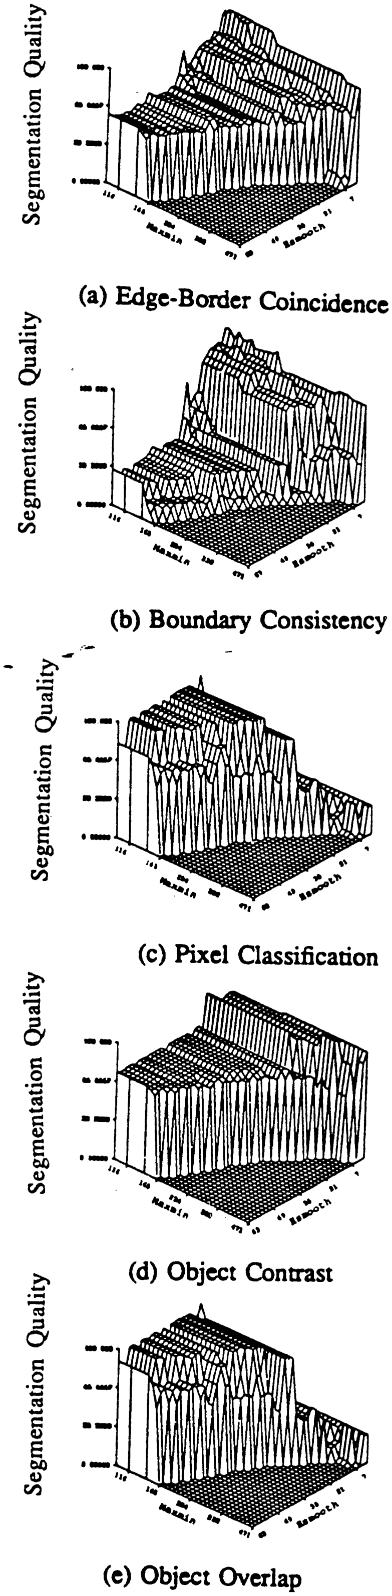

(a) Edge-Border Coincidence

(b) Boundary Consistency

(c) Pixel Classification

(d) Object Contrast

(e) Object Overlap

Figure 4: Individual quality surfaces

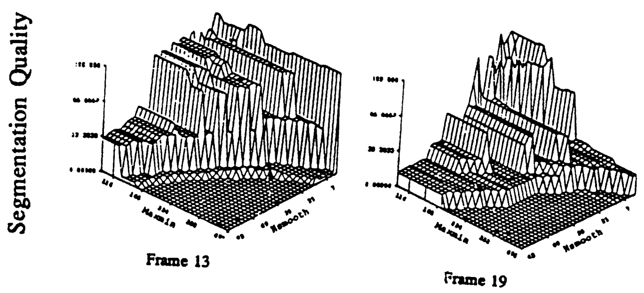

Figure 6: Quality Surfaces for the images shown in Fig. 5

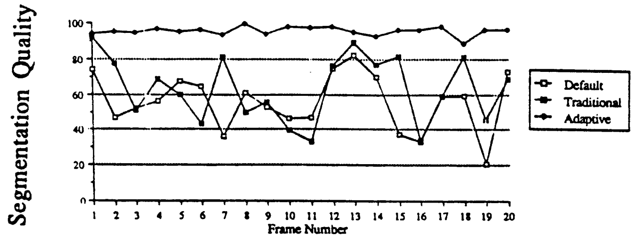

Figure 7: Comparison of the adaptive image segmentation system with default Phoenix parameters and the traditional segmentation approach commonly used in the computer vision field.

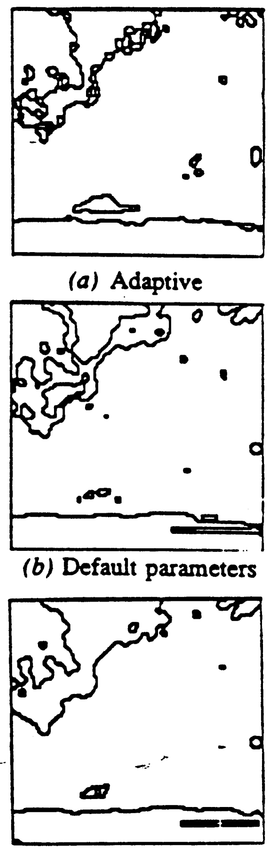

(a) Adaptive

(b) Default parameters

(*c*) Traditional Approach

Figure 8: Segmentation performance comparison.

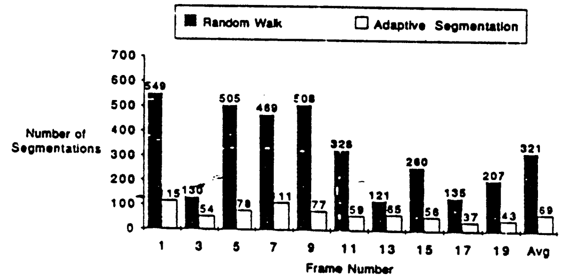

Figure 9: Performance comparison of the adaptive segmentation technique and a random walk approach.

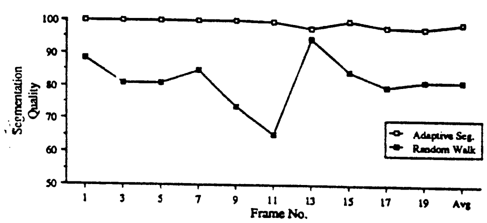

Figure 10: Segmentation quality performance for the adaptive segmentation technique and the random walk approach.

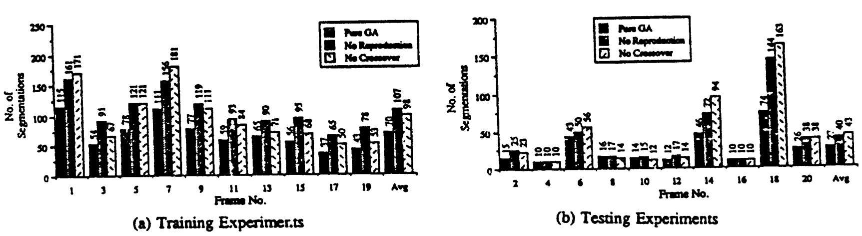

Figure 11 Performance comparison of the pure genetic algorithm and its variations. The superior performance of the pure genetic algorithm demonstrates the effectiveness of the reproduction and crossover operators.

PROCEEDINGS OF THE FOURTH INTERNATIONAL CONFERENCE ON

GENETIC ALGORITHMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

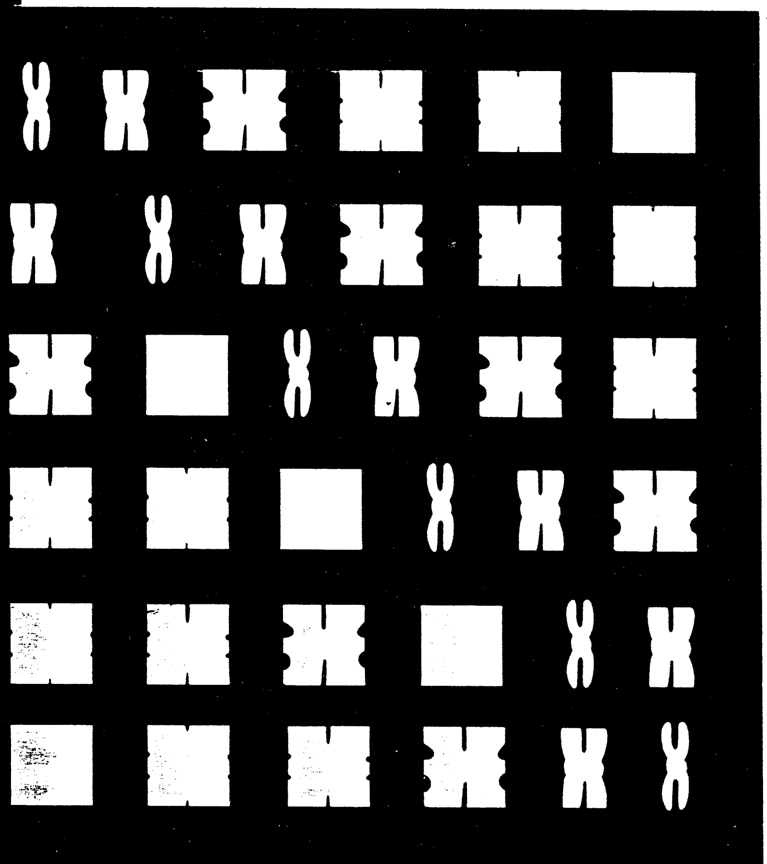

RICHARD K. BELEW
UNIVERSITY OF CALIFORNIA, SAN DIEGO

LASHON B. BOOKER
MITRE CORPORATION

EDITORS

PROCEEDINGS OF THE
FOURTH INTERNATIONAL
CONFERENCE ON

GENETIC ALGORITHMS

University of California, San Diego
July 13–16, 1991

Editors/Program Co-Chairpersons:
Richard K. Belew
Lashon B. Booker

Supported By:
International Society for Genetic Algorithms
Office of Naval Research
Naval Research Laboratory

Morgan Kaufmann Publishers
San Mateo, California

Sponsoring Editor Bruce Spatz
Production Editor Yonie Overton
Cover Design Jo Jackson
Production Coordination Ocean View Technical Publications

Morgan Kaufmann Publishers, Inc.
Editorial Office:
2929 Campus Drive, Suite 260
San Mateo, CA 94403

94 93 92 91 5 4 3 2 1

Library of Congress Cataloging in Publication Data is available for this book.
ISBN 1-55860-208-9