A Platform for Evolving Genetic Automata for Text Segmentation (GNATS)

Michael D. Garris
National Institute of Standards and Technology,
Gaithersburg, Maryland 20899

Proceedings: SPIE Vol. 1710 Science of Artificial Neural Networks
Orlando, April 1992

ABSTRACT

Developers of large-scale document processing and image recognition systems are in need of a dynamically robust character segmentation component. Without this essential module, potential turn-key products will remain in the laboratory indefinitely. An experiment of evolving a biologically-based neural image processing system which has the ability to isolate characters within an unstructured text image is presented. In this study, organisms are simulated using a genetic algorithm with the goal of learning the intelligent behavior required for locating and consuming text image characters. Each artificial life-form is defined by a genotype containing a list of interdependent control parameters which contribute to specific functions of the organism. Control functions include vision, consumption, and movement. Using asexual reproduction in conjunction with random mutation, a domain independent solution for text segmentation is sought. For this experiment, an organism's vision system utilizes a rectangular receptor field with signals accumulated using Gabor functions. The optimal subset of Gabor kernel functions for conducting character segmentation are determined through the process of evolution. From the results, two analyses are presented. A study of performance over evolved generations shows that qualifiers for the natural selection of dominant organisms increased 62%. The second analysis visually compares and discusses the variations of dominant genotypes from the first generation to the uniform genotypes resulting from the final generation.

1. INTRODUCTION

Research in neural classification techniques have shown great potential for improving the accuracy of character recognition engines. Character recognition, the classification of well formed and cleanly segmented characters, has been studied in great detail in the past.[1-5] What is often avoided in character recognition research is the study of automated segmentation, the separation of text images into individual letters, one letter per image. The success of character-based classifiers is directly dependent upon successful segmentation of text images into isolated characters. Full scale document recognition and processing systems must include an accurate / intelligent segmentation component.

A model recognition system designed to process handprint written on structured forms has been implemented on a massively parallel computer at NIST.[6] This model system has been used to study the effect of segmentation methods on recognition accuracy and system throughput. Conventional approaches to character segmentation involve algorithms based on spatial histograms, connected-component labeling, and heuristics for resolving ambiguities.[7] The rules used in conventional approaches depend upon the intrinsic structure of a specific text, the size and style of machine print, and whether or not the text image contains handwriting. Studies have shown that methods utilizing traditional image processing techniques such as spatial histograms are only 60% accurate on handprint, and even when implemented on a parallel computer, require 55% of the system's processing time.[6,7] In order to develop generalized segmentation solutions which are both accurate and efficient, alternative methods are being explored.

2. GENETIC ALGORITHMS WITH CONNECTION-ORIENTED MODELS

The experiment presented in this paper explores an artificial life-form which exhibits the necessary behavior to seek out and isolate characters, domain independently. The integration of genetic algorithms and evolution-based learning with connection-oriented models shows great promise for application solutions in the future. Vasant Honavar and Leonard Uhr have introduced the notion of "Generalized Connectionist Networks."[8] In their discussion of generalized network solutions, they define a set of functional subsystems which can apply genetic search and evolution-based learning to build connection topologies, find sets of control parameters, and select learning strategies which provide optimal solutions.

Experiments involving the integration of evolution with neural learning have been conducted, including work at the Center for Research in Language at the University of California, San Diego.[9,10] In these experiments, an organism was simulated which, through the use of natural selection, reproduction, and mutation, exhibited increasingly intelligent behavior. In this case, the desired behavior was for an organism to strategically move through its grid-oriented environment finding and eating cells of food.

The experiment presented in this paper is similar in that an organism is to be simulated which exhibits intelligent behavior for locating and consuming text image characters. A domain independent solution for text segmentation is sought by evolving Genetic Neural Automata for Text Segmentation (GNATS). Using evolution-based techniques for learning, the potential for solving an engineering application with biologically motivated methods is demonstrated.

3. GENETIC MODEL

A GNAT is fundamentally defined and represented by a genotype. Each control parameter can therefore be thought of as a gene which contributes to a specific function of the organism and which is interrelated with other genes within the genotype. Control functions include vision, consumption, and movement. These genotypes are passed from generation to generation through a process of natural selection and reproduction including random mutations from parent to child. By monitoring the success of each GNAT within its lifetime, only the fittest are selected for reproduction. A child GNAT therefore has the potential for performing as well as its parent, performing better than its parent, or failing to perform altogether. Goal directed behavior, in theory, should therefore increase as GNATS are selected and reproduce from generation to generation.

3.1 Gabor-based Vision System

In order for a GNAT to sense its environment, a vision system and sensing mechanism is required. For this experiment, the vision system utilizes a rectangular receptor field which can be used to scan the organism's immediate surroundings. The efficiency of this vision system will contribute directly to the success or failure of the organism.

Within the receptor field, signals are accumulated using a set of incomplete nonlinear functions, the Gabor functions. This receptor field model is based on known biological data on the structure of vertebrate vision but is implemented using more conventional numerical methods.[11] Gabor functions reduce random image noise and smooth irregularities in image structure by acting as spatially localized low-pass filters. John Daugman[12] has used Gabor functions for image compression and image texture analysis. These functions match the visual receptor field profiles of mammalian eyes and provide the minimum combination of uncertainty in position and spatial frequency.

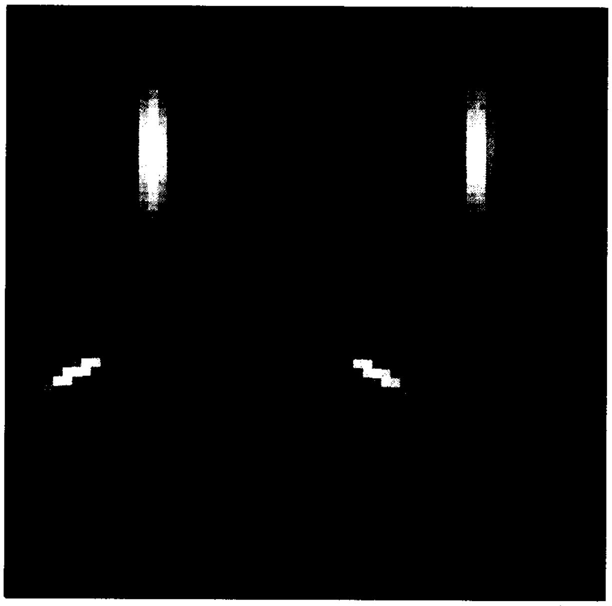

Figure 1. Four different Gabor kernel functions.

Figure 1 displays several Gabor functions demonstrating the ability to select functions based on position, orientation, and spatial frequency. The image displayed in the top left quadrant is an example of a Gabor function defined with cosine phase. The image in the top right quadrant is a similar function defined with sine phase. Both have the same spatial extent and are vertically oriented. The image displayed in the bottom left quadrant of the figure is an example of a Gabor function defined one level down from the one displayed immediately above it. Levels are discussed later in this section. Notice that this function is half the spatial extent, both horizontally and vertically, and has been oriented at 30 degrees. The bottom right Gabor function is similar to the one displayed on the bottom left, only it is being displayed at an orientation of 150 degrees.

Since the Gabor basis functions are an infinite set, it is necessary to select a specific subset to be used as the image reconstruction elements which cover the character image. This selection process is referred to as tiling and is defined by a set of attributes which include horizontal and vertical starting levels, the number of extending levels, the number of evenly spaced orientations, and function symmetry.[11] First level functions by definition have a spatial extent the size of the image being reconstructed and share a common origin at the center of the image. Successive levels represent recursive subdivisions of the original image space specifying tiles with smaller and smaller spatial extent and with origins at twice the sampling density of the previous level.

Figure 2 illustrates the result of tiling a square image starting at horizontal and vertical levels equal to 0 and extending 2 levels. Five different tiles are created. The first level contains one tile the size of the entire image, d_0 by d_0, while the second level contains four tiles, each of size d_1 by d_1, which are 1/4 the area of the tile in level 1. Five different function origins are shown in Figure 2, each centered in the middle of their corresponding tile. Given this pattern of recursive subdivision, level 3 would contain 16 tiles. Unlike this example, the horizontal and vertical starting levels may differ causing spatial subdivisions to be asymmetrically proportional. In this way, assigning values for the starting levels and the number of extending levels dictates the number of resulting tiles along with their spatial extent and origin.

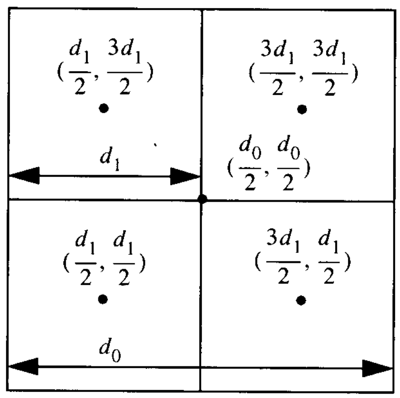

Figure 2. Recursive subdivision of image space used in Gabor function tiling.

At each tile origin, there may exist one or more Gabor functions rotated at even intervals according to the number of orientations specified. For example in Figure 2, an orientation value of 4 would define 20 different (origin, orientation) pairs, four different orientations at each of the five tile origins. The four orientations within each tile would be evenly distributed at intervals of 45 degrees, 180 degrees divided by 4. The range of possible orientations, 360 degrees, is reduced by a factor of two in order to eliminate orientation redundancies. The symmetry attribute determines the phase of the selected functions and may specify a single set of Gabor functions with either sine or cosine phase. Alternatively, the symmetry attribute may be used to specify both sine and cosine phases resulting in two sets of Gabor functions selected for each (origin, orientation) pair. In our example, specifying either sine *or* cosine symmetry would result in 20 selected Gabor functions, whereas specifying both sine *and* cosine symmetry would result in 40 selected Gabor functions.

The optimal subset of Gabor kernel functions for conducting character segmentation will be determined by evolving values for the tiling attributes described above. In this way, multiple topologies of the vision system can be explored and evaluated across successive generations of GNATS. Using a selected subset of incomplete nonlinear Gabor functions, an image is reconstructed using a least squares fit. The coefficients resulting from the fit can be accumulated to produce a broad signal for detecting the presence of significant image information within the view of the receptor field.

3.2 Consumption

The reconstructed image resulting from the least squares fit is used to define the boundaries within which the GNAT is to consume a sensed character. The grayscale reconstructed image is thresholded to binary and then used as a logical mask for extracting (segmenting) the information (characters) from the GNAT's environment. The two images displayed in Figure 3 demonstrate the ability to isolate words by exploiting the Gabor functions' body detecting capability. Similar results exist for individual characters. Notice how the individual letters of the digitized word "Constitution" are collapsed into a single mass of black pixels in the Gabor reconstructed image.

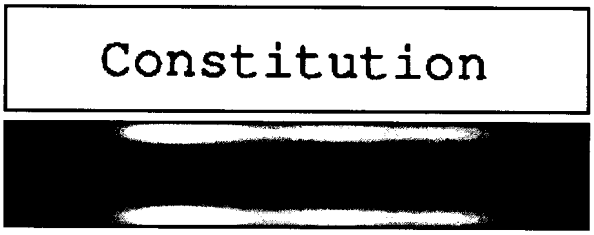

Figure 3. Example of a Gabor image reconstruction.

3.3 Movement

Primitive motion capabilities are provided to give a GNAT the ability to explore its environment. A GNAT's neighborhood is defined as its eight immediately surrounding receptor field views. If a GNAT senses significant image information within its neighborhood, then the organism moves a fraction of its receptor field size in the direction of the sensed signal. If no significant signal is sensed, then the GNAT jumps in a random direction just beyond its current neighborhood.

3.4 Environment

An external control is required in order to govern how a GNAT is to actually live and grow in its environment. It is within the environment that the application domain is modeled and success measured. GNATS are placed into isolated environments, one GNAT per environment. A GNAT, over its lifetime, will have opportunity to explore a predetermined number of consistent and yet unique environments; consistent in complexity, yet unique in layout. Each environment is defined as a rectangular binary image containing one or more words randomly distributed throughout the image space. Words found in an environment are synthesized from a database of digitized characters so that the location and content of each word is registered for future reference. Word synthesis can be thought of as an image-based type setter.

GNATS are required to move about within their environments sensing and locating isolated character images and extracting them from the environment using the control components described above. A GNAT is permitted to explore a given environment for a predetermined number of computer clock cycles during which time a qualitative value representing the number of consumed characters is accumulated.

4. EXPERIMENT: EVOLVING PARAMETERIZED CONTROL SETTINGS

This section lists the internal details of the experiment conducted which was based upon the ideas presented above. For the purpose of proving feasibility and avoiding unnecessary complexities, this initial study was designed to be complete in terms of theory and yet simplistic in terms of solution.

4.1 Code Ranges for Genotype Initialization and Mutation

Evolution-based learning provides a platform whereby optimal sets of internal GNAT control parameters can be determined. To initiate this experiment, ranges for each control parameter in the genotype structure were assigned. The limits set by the ranges are used as guidelines controlling the random initialization of first-generation GNATS and the magnitude of valid mutations occurring during reproduction of generations thereafter. The genotype codes and their associated ranges are listed in Table 1.

PARAMETER	MIN	MAX
GNAT_MOUTH_WIDTH	30	256
GNAT_MOUTH_HEIGHT	30	90
GBR_X_LEVEL	0	3
GBR_Y_LEVEL	0	3
GBR_LEVELS	1	2
GBR_SYMMETRY	-1	1
GBR_THETAS	1	2
GBR_THRESH	64	192
GNAT_SIGNAL_THRESH	0	100
GNAT_DELTA	1	10

Table 1. Range values for genotype initialization and mutation.

Given these minimum and maximum limits, the width of a GNAT's receptor field is permitted to mutate between 30 and 256 pixels. This range corresponds approximately to field widths varying from 1 to 8 characters in size when using machine printed characters digitized at 300 dpi for environment synthesis. The height of the receptor field is permitted to mutate allowing field heights to vary between approximately 1 to 3 characters in size. Gabor function tiling is permitted to begin at levels 0 to 3 independently along both the X and Y axes and will span a maximum of 2 levels from the selected starting levels. The symmetry of the tiled functions may include sine and / or cosine frequency components. The number of orientations for each tiled function may be horizontal or both horizontal and vertical. Through reproductive mutations across these dynamic ranges, an optimal set of vision parameters can be determined across evolved generations.

The consumption threshold used to extract the sensed image information from within a GNAT's grayscale receptor field ranges between 64 and 192. The focused range therefore centers the potential threshold values between 25% to 75% of the full dynamic range of possible 8-bit pixel values. The signal threshold for sensing significant image information within a GNAT's immediate receptor field neighborhood is permitted to mutate between 0 to 100. As the selected threshold value increases, a GNAT will be more and more selective to what it moves towards. If significant information is sensed, then the GNAT will move a fraction of the receptor field size in the direction of the sensed information. This fractional increment is permitted to mutate in integer multiples of tenths of the receptor field size.

4.2 Generation Control Parameters

In addition to genotype ranges, global architecture parameters for environment synthesis and generation control are also necessary. Table 2 lists the parameters used to control the evolution of GNATS from one generation to the next.

PARAMETER	VALUE
GNAT_POP	100
GNAT_GENERATION	15
GNAT_SURVIVE	10
GNAT_CHILDREN	10
GNAT_MUTATE	2

Table 2. Parameter values for generation control.

This experiment maintained a stable generation population of 100 GNATS and monitored the performance of 15 successive GNAT generations. From each simulated generation, the top 10 dominating GNATS were allowed to reproduce. The accumulated qualifier used for natural selection in this experiment was the total number of black pixels consumed across the lifetime of each GNAT. The GNATS who consumed the greatest quantity of pixels within their lifetime were permitted to replicate their genotypes. Each of the 10 dominant GNATS was permitted to have 10 children thus maintaining the stable 100 GNAT population across generations. In this initial experiment, asexual reproduction was conducted so that each child received a replicated genotype from a single parent with a small amount of mutation. In each inherited genotype, 2 randomly chosen codes out of the 10 possible codes (20%) were altered, and those varied within the ranges dictated in the Table 1.

4.3 Environment Synthesis and Control Parameters

Table 3 lists the parameter values used to synthesize and initialize the environments used in this experiment. GNATS were permitted to explore five unique but consistently complex environments. The width and height of each environment was 608 by 608 pixels containing 20 registered words. For initial simplicity, all words were the single character word "a" randomly distributed through the environment space. Every GNAT across the 15 generations of this experiment explored the same 5 environments, each in isolation. An image of one of the 5 environments used in the experiment is displayed in Figure 4.

PARAMETER	VALUE
ENV_NUM	5
ENV_WIDTH	608
ENV_HEIGHT	608
ENV_WORD_COUNT	20
ENV_WORD	"a"
ENV_LSPAN	7200
ENV_STRT_X	300
ENV_STRT_Y	300

Table 3. Parameter values for environment synthesis and control.

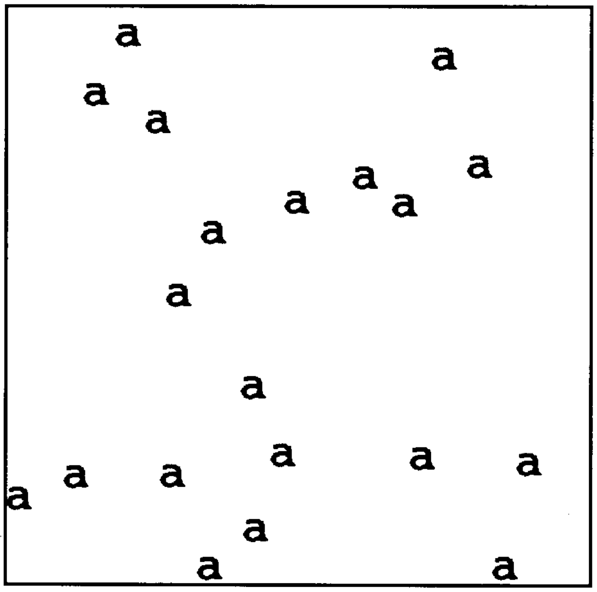

Figure 4. One of 5 environments used in the experiment.

A GNAT was permitted to live 7,200 CPU clock cycles, 2 minutes, within each of its 5 environments totalling an overall life-span of 10 minutes of CPU time. Due to GNATS having equal life-spans, each one had equal potential for exhibiting successful behavior. Simulating a generation of 100 GNATS required no less than 1,000 CPU minutes. This experiment ran at a low priority for approximately 20 days on a scientific workstation with only minimal interruptions. Each GNAT was initially placed in the center of each environment at pixel coordinate (300, 300) and then permitted to explore based on its ability to sense and move.

5. RESULTS

This section discusses the results of conducting the experiment outlined in Section 4. First, the performance of the GNATS evolving over time is examined, and then an analysis of how well evolution-based techniques promoted the convergence of genotype codes across generations is presented.

5.1 Behavior Across Generations

The accumulated qualifier used for natural selection in this experiment was the total number of black pixels (text pixels) consumed over a GNAT's lifetime. Figure 5 plots the consumed pixels for the top 10 performers from each generation. The solid line charts the average number of accumulated pixels across the 10 GNATS whereas the dotted line charts the maximum number of black pixels eaten by a single GNAT within each generation. In both instances, the overall performance of the dominant GNATS over generations increased leveling off about the 8th generation.

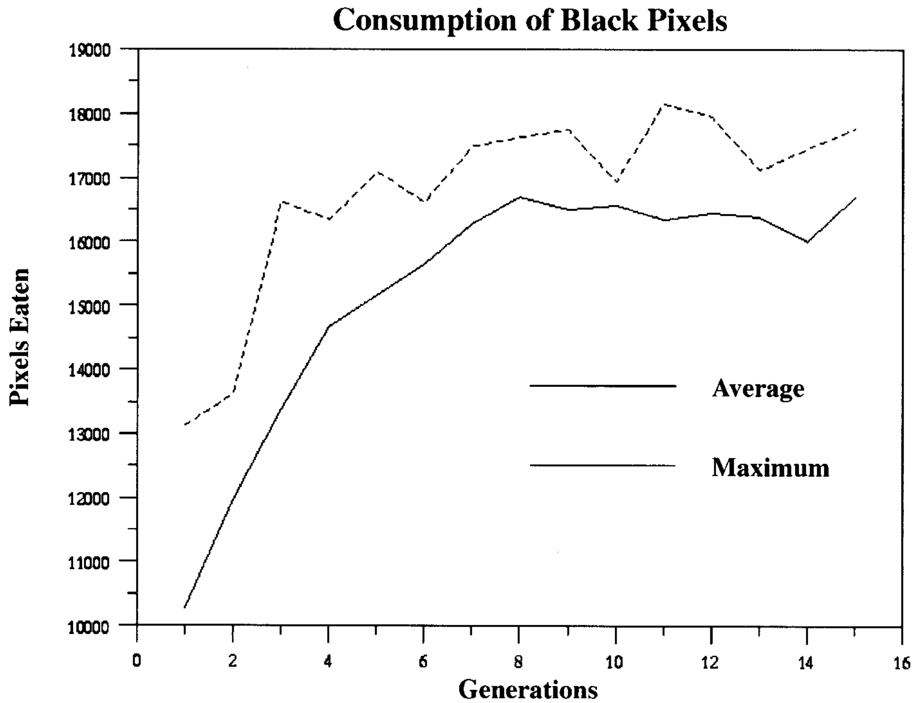

Figure 5. Graph representing the amount of pixels consumed by dominant GNATS over evolved generations.

Table 4 lists the genotype possessed by the GNAT which consumed the greatest number of pixels over all 15 generations. This GNAT was a member of generation 10 and consumed 18,163 black pixels over its lifetime. Studying the individual code values shows that maximum consumption was achieved through a relatively simplistic control structure. This demonstrates that efficient control parameters can be derived from a set of potentially complex solutions through the techniques of evolution. This GNAT has a minimal size receptor field, a computationally inexpensive vision system, and a very low signal threshold which causes it to be non-selective about sensed image information.

PARAMETER	VALUE
GNAT_MOUTH_WIDTH	40
GNAT_MOUTH_HEIGHT	31
GBR_X_LEVEL	0
GBR_Y_LEVEL	0
GBR_LEVELS	1
GBR_SYMMETRY	0
GBR_THETAS	2
GBR_THRESH	180
GNAT_SIGNAL_THRESH	1
GNAT_DELTA	5

Table 4. Genotype of the GNAT which consumed the greatest number of pixels across all generations.

The first 20 bites taken by the GNAT described in Table 4 are displayed in Figure 6. Notice how the small receptor fields have fragmented characters into pieces. This GNAT exhibits extremely primitive behavior; it has a vigorous appetite and is non-selective about what it consumes. Therefore, this GNAT consumed the greatest number of text pixels, but does not exhibit the originally desired behavior of isolating complete characters of text.

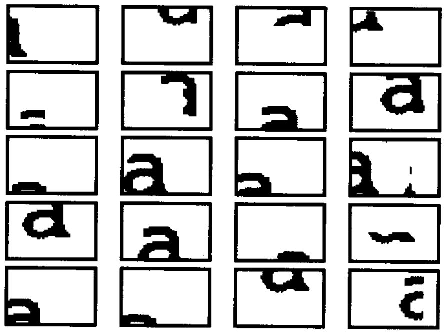

Figure 6. First 15 bites taken by the GNAT which consumed the greatest number of pixels across all generations.

Compare the GNAT described above with the GNAT whose genotype is listed in the Table 5. This GNAT consumed 16,947 black pixels, which was the maximum within generation 9. Its code values are quite similar to the GNAT in Table 4 with only the size of the receptor field differing significantly. The first 15 bites of this GNAT are displayed in Figure 7. These images illustrate a more desired behavior than those bites shown in Figure 6. The bites made by the second GNAT contain a greater percentage of whole characters minimizing the amount of fragmented character pieces.

PARAMETER	VALUE
GNAT_MOUTH_WIDTH	40
GNAT_MOUTH_HEIGHT	72
GBR_X_LEVEL	0
GBR_Y_LEVEL	0
GBR_LEVELS	1
GBR_SYMMETRY	-1
GBR_THETAS	2
GBR_THRESH	180
GNAT_SIGNAL_THRESH	2
GNAT_DELTA	5

Table 5. Genotype of the GNAT which consumed the greatest number of pixels in generation 9.

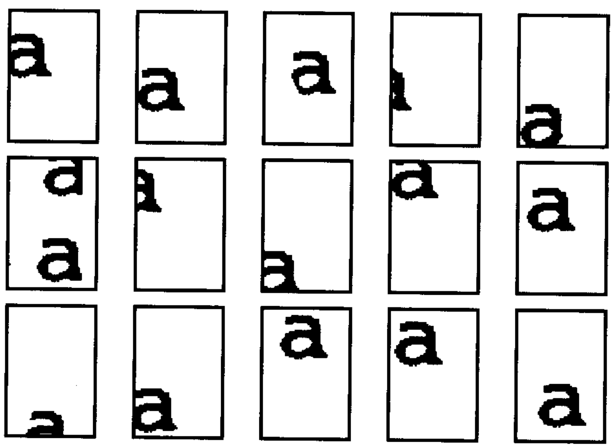

Figure 7. First 15 bites taken by the GNAT which consumed the greatest number of pixels in generation 9.

The behavior of the GNAT in Table 5 did not perpetuate itself through successive generations because of two primary factors. First, as the size of a receptor field increases, the computational expense of the Gabor-based vision system also increases. Therefore, a vision system such as the one defined in Table 4 will have the opportunity, within a GNAT's lifetime, to take many more bites than one which is more computationally complex. This type of competition is required to derive solutions which are not only effective, but also efficient. The second factor deals with the accumulated qualifier used for natural selection. In this experiment, this qualifier was the number of consumed black pixels. Notice that this criterion contains no intrinsic rewards for bites which consume whole characters nor are there intrinsic penalties for bites which consume partial characters. This over-simplified qualifier in conjunction with a small receptor field will therefore win the right of reproduction over a GNAT such as the one in Table 5.

5.2 Analysis of Evolved Genotypes

In this section, the convergence of genotype codes across generations is analyzed. The graph in Figure 8 contains a plot of 9 normalized genotypes belonging to the 9 GNATS from generation 1 which consumed the greatest number of pixels. The unit points, 1 through 10, along the X-axis correspond in order to the genotype codes listed in Table 1. The values plotted along the Y-axis represent the normalized values associated with each of the codes. Code values were normalized in the range 0.0 to 1.0 according to the parameter limits listed in Table 1. The graph in Figure 8 is a 2-dimensional representation allowing visual inspection of the 10-dimensional genotype.

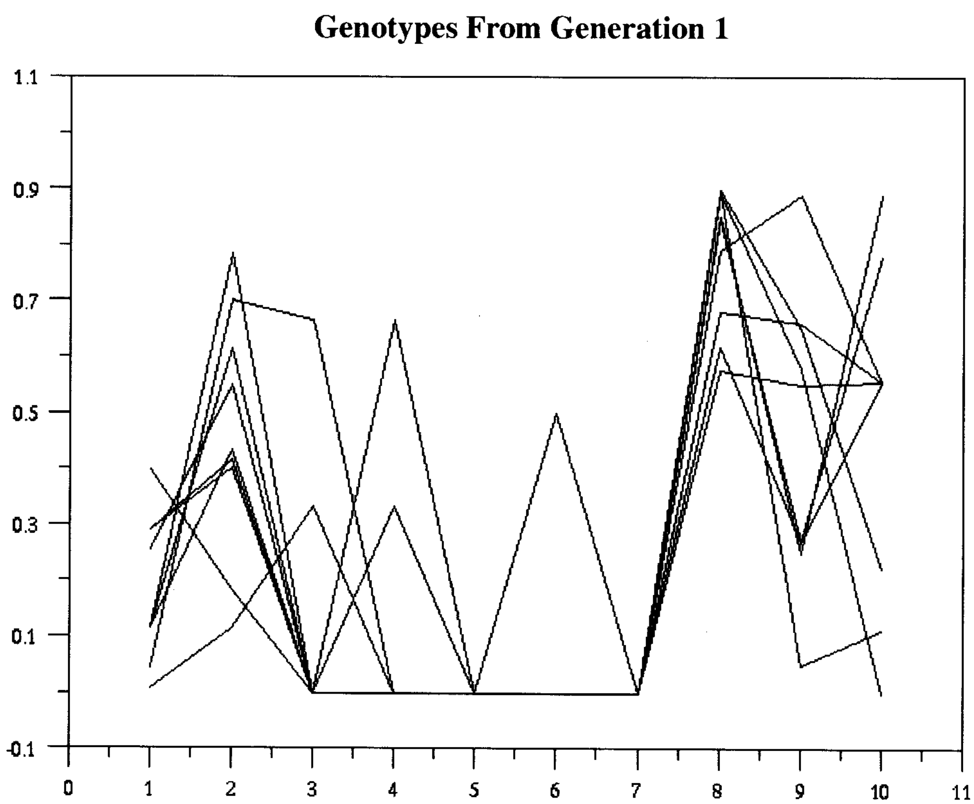

Figure 8. Graph of the top 9 normalized genotypes from generation 1.

Compare the large variations among the genotypes in the graph above to the genotypes plotted in Figure 9. The second plot shows the normalized genotypes belonging to the 9 GNATS in generation 15 which consumed the greatest number of pixels. These genotypes are much more uniform, demonstrating that the genetic search converged to a small number of solution sets from what was originally an extremely large number of potential solutions. A breakdown of the genotypes displayed in Figure 9 is given below.

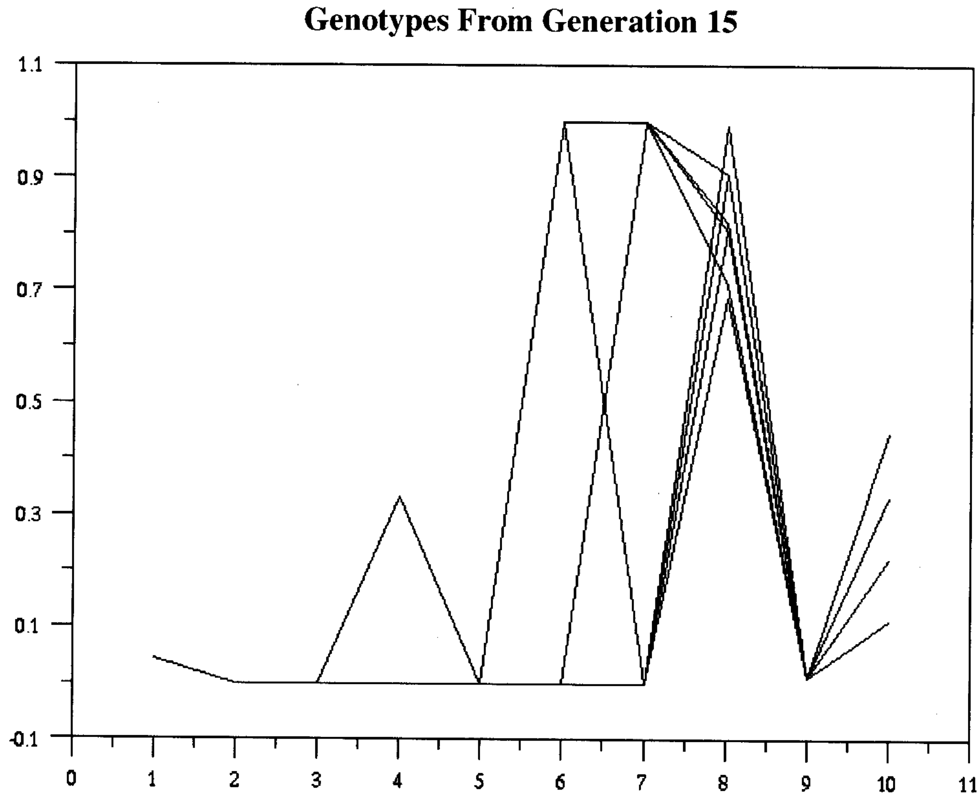

Figure 9. Graph of the top 9 normalized genotypes from generation 15.

Upon simulation of generation 15, the size of the 9 receptor fields represented by codes 1 and 2 in Figure 9 are minimal in size similar to the GNAT described in Table 4. Codes 3, 4, and 5 correspond to the tiling levels used in the Gabor-based vision system. All 9 GNATS contained the same 3 parameter levels with the exception of one GNAT which deviated on code 4. The symmetry switch for selecting frequency components of the Gabor functions is represented by code 6. The dominant GNATS chose one component or the other with none choosing both. The 9 GNATS split on values for code 7 with 5 GNATS using horizontal orientations and 4 GNATS using both horizontal and vertical orientations of their selected Gabor functions. All the GNATS had a consumption threshold which was relatively large represented by the values plotted for code 8. Code 9, the signal threshold, is consistently very low implying that the GNATS were non-selective. The last code in the right graph represents the percentage of receptor field size used to move in the direction of sensed information. This code remains below 50% of the receptor field size. In light of these observations, this analysis clearly demonstrates the converging power of genetic search.

6. OBSERVATIONS AND CONCLUSIONS

This experiment examines the potential for using genetic algorithms and artificial life-forms for segmenting images of text. By design, the initial study was simplistic therefore limiting the usefulness of the solution sets derived. None the less, much can be gained from the results.

Two general observations can be made from the analyses in Section 5. First, the extent to which desired behavior is evolved is directly dependent upon the criterion used for natural selection. In the case of this study, the accumulated number of eaten black pixels did not reward GNATS which ate whole characters nor did it penalize GNATS which ate pieces of characters. This allowed extremely efficient but primitive GNATS to dominate more sophisticated GNATS which exhibited more desirable behavior.

The second observation deals with the problem of predeterminism. Unlike traditional artificial neural models, the genetic model presented here evolved the parameters governing the control of the organism. This allowed the control structure to be dynamically derived. However, predeterminism is not completely avoided. If the codes within the genotype are not designed properly, then the organisms being evolved will be fundamentally handicapped.

In conclusion, this study has successfully shown that genetic algorithms can be used to derive efficient solution sets from a large collection of complex and dynamic control parameters. The results from the initial experiment demonstrate the feasibility of evolving an artificial life-form capable of segmenting text images. The experiment also demonstrates great potential for further research and development. Future work will include the refinement of the natural selection criterion and the integration of neural learning for dynamic sensing and focusing into the genetic platform established here.

7. REFERENCES

1. M. D. Garris, R. A. Wilkinson, and C. L. Wilson, "Methods for Enhancing Neural Network Handwritten Character Recognition," *International Joint Conference on Neural Networks*, Vol. I, pp. 695-700, Seattle, 1991.
2. G. E. Hinton, and C. K. I. Williams, "Adaptive Elastic Models for Character Recognition," *Advances in Neural Information Processing Systems*, R. Lippmann, Vol. IV, to be published, Denver, 1991.
3. L. D. Jackel, H. P. Graf, W. Hubbard, J. S. Densker, D. Henderson, and Isabelle Guyon, "An Application of Neural Net Chips: Handwritten Digit Recognition," *International Joint Conference on Neural Networks*, Vol. II, pp. 107-115, San Diego, 1988.
4. G. L. Martin and J. A. Pittman, "Recognizing Hand-Printed Letters and Digits," *Advances in Neural Information Processing Systems*, D. S. Touretzky, Vol. 2, 405-414, Morgan Kaufmann, Denver, 1989.
5. C. L. Wilson, R. A. Wilkinson, and M. D. Garris, "Self-Organizing Neural Network Character Recognition on a Massively Parallel Computer," *International Joint Conference on Neural Networks*, Vol. II, pp. 325-329, San Diego, 1988.
6. M. D. Garris, C. L. Wilson, J. L. Blue, G. T. Candela, P. Grother, S. Janet, and R. A. Wilkinson, "Massively Parallel Implementation of Character Recognition Systems," SPIE Machine Vision Applications in Character Recognition and Industrial Inspection, to be published, San Jose, 1992.
7. R. A. Wilkinson, "Segmenting Text Images with Massively Parallel Machines," SPIE Intelligent Robots and computer vision X, to be published, Boston, 1991.
8. Vasant Honavar and Leonard Uhr, *Symbol Processing Systems, Connectionist Networks, and Generalized Connectionist Networks*, Technical Report #90-23, Department of Computer Science, Iowa State University, December 1990.
9. Stefano Nolfi, Jeffrey L. Elman, and Domenico Parisi, *Learning and Evolution in Neural Networks*, CRL Technical Report 9019, Center for Research in Language, University of California, San Diego, July 1990.
10. Valentino Braitenberg, *Vehicles, Experiments in Synthetic Psychology*, MIT Press, Cambridge, Massachusetts, 1984.
11. C. L. Wilson, R. A. Wilkinson, and M. D. Garris, "Self-Organizing Neural Network Character Recognition Using Adaptive Filtering and Feature Extraction," to be published in *Progress in Neural Networks*, Volume 3, 1991.
12. J. G. Daugman, "Complete Discrete 2-D Gabor Transform by Neural Networks for Image Analysis and Compression," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, **36**, pp. 1169-1179, 1988.