Evaluating Mutation Operators for Evolved Image Reconstruction Transforms

Michael R. Peterson
Dept. of Computer Science and Engineering
Wright State University
Dayton, OH, USA
peterson.7@wright.edu

Gary B. Lamont
Dept. of Electrical and Computer Engineering
U. S. Air Force Institute of Technology
WPAFB, OH, USA
gary.lamont@afit.edu

Frank Moore
Dept. of Mathematical Sciences
University of Alaska Anchorage
Anchorage, AK, USA
affwm@uaa.alaska.edu

ABSTRACT

Various military systems require image and signal processing, often in noisy or bandwidth-limited situations. In this research, we employ genetic algorithms (GAs) to evolve forward and inverse transforms that reduce quantization error in reconstructed signals and images. The resulting transforms produce higher quality images than current wavelet-based transforms at a given compression ratio and thus allow transmission of compressed data at a lower bandwidth. We expand on previous research by evaluating several mutation strategies for evolving reconstruction filters. Our results indicate that GAs employing Gaussian mutation applied with shrinking standard deviations evolve transforms superior to transforms evolved by GAs employing other tested mutation operators.

Categories and Subject Descriptors

J.2 [**Physical Sciences and Engineering**]: Engineering

General Terms

Algorithms, Performance, Experimentation

Keywords

Genetic algorithms, wavelets, optimization, mutation

1. INTRODUCTION

Image and signal processing are active areas of military research. Satellites and Unmanned Aerial Vehicles (UAVs) potentially collect huge amounts of image data during surveillance missions. Sonar and radar systems process huge amounts of sensor data in real time. The requirements to minimize mission cost while maximizing effectiveness necessitates the development of compression techniques that simultaneously minimize storage and bandwidth requirements while maintaining maximum signal information.

With these requirements in mind, quantization of digitized data is often necessary for military DSP applications. *Quantization* minimizes storage requirements by mapping all values in signal x to a small discrete range of values $Q(x)$. Though quantization greatly improves compression ratios, perfect reconstruction of x from $Q(x)$ is impossible due to the loss of low-order bits [8]. *Wavelets* [1] are a standard methodology for signal compression algorithms. The discrete wavelet transform (DWT) redistributes the energy in a signal by transforming a time signal into a time-frequency domain. A signal may be compressed by first applying the DWT, followed by quantization, and then by applying entropy coding. Signals are reconstructed in a reverse manner. Most information loss occurs during quantization.

Figure 1 shows a publicly available satellite image of the U.S. Air Force Museum in Dayton, OH, taken approximately in 2003 [3]. This image shows a diverse set of aircraft, two existing hangers, and a third hanger under construction. This low-resolution image is representative of the type of image that may be obtained by an expendable UAV during a reconnaissance mission in hostile territory. Figure 2 shows the same image after it has been compressed with the Daubechies-4 (DB4) DWT, quantized with a quantization step of 64, dequantized, and reconstructed by the DB4 inverse transform. Note the loss of information due to quantization. A military grade system must minimize this loss to maximize the intelligence that may be gathered while maintaining a desired compression ratio. Our previous research established that GAs are capable of evolving coefficient sets defining filters that provide improved signal and image reconstruction in comparison to the DWT under conditions subject to quantization error [5, 7]. We seek to improve GA performance by evaluating several real-valued mutation operators.

2. METHODOLOGY

We compare the performance of GAs using four different mutation operators designed for real-valued operators to evolve image reconstruction filters. The first operator employs mutation steps sampled from a Gaussian mutation centered at the current gene value with a standard deviation that shrinks by generation [2]. The second operator is related to the first, but new values are instead sampled from a Cauchy distribution. This distribution has longer tails than

Figure 1: Satellite image of US Air Force Museum.

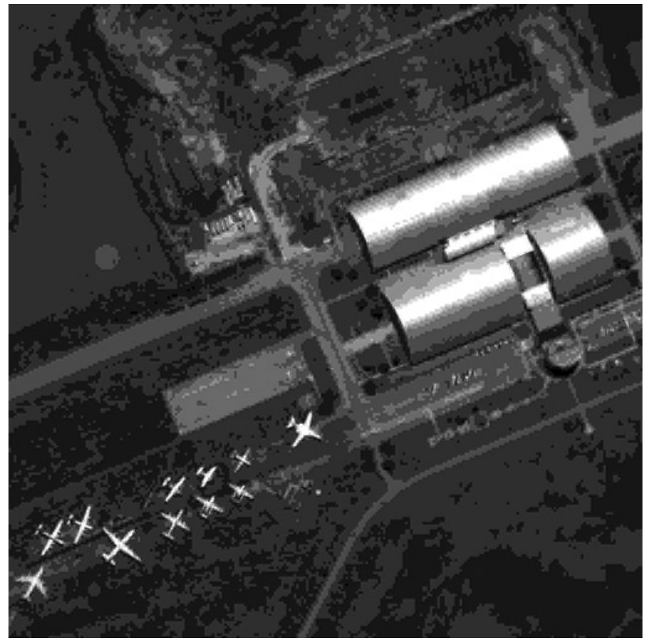

Figure 2: Reconstructed image after quantization.

the Gaussian distribution and hence samples a wider area of the search space [10]. The third operator is a nonuniform operator originally proposed by Michalewicz [4]. Genes are mutated according to distributions that have ranges shrinking by generation and that favor a balanced search between the minimum and maximum bounds on the genes. The final mutation operator [7] applies a small random step change to the current gene value for 95% of all mutations, and flips the sign of a gene in 5% of the mutations. We refer to this mutation as a local mutation.

The GA evolves a population of 50 individuals for 500 generations in each experiment. Parents are chosen through stochastic uniform selection. Recombination occurs through heuristic crossover. A set percentage of children are created through crossover, the remaining are created by mutation. This crossover rate is empirically set to 0.7, 0.35, 0.65, and 0.7 for Gaussian, Cauchy, nonuniform, and local mutation, respectively. The two most fit parents survive to the next generation, and the remaining population consists of children built through crossover and mutation. The universal quality index (UQI) measuring image quality [9] defines the fitness function. UQI may range from -1 to 1, with a value of 1 indicating perfect reconstruction. Experiments are conducted at one level of DWT decomposition at quantization level 64 on the standard greyscale fruits test image [7], chosen for its varieties of pixel intensities and textures.

20 experiments are conducted for each mutation operator, using two population initialization routines. The first initializes individuals randomly, sampling a uniform distribution. This technique tests the GA's ability to sample a global search space. The second routine creates initial individuals located near the DB4 reconstruction coefficients. This technique tests the GAs' abilities to refine wavelet coefficients.

Random Population Initialization Fitness

mutation	NonUniform	Cauchy	Local	Gaussian
	Descriptive Statistics			
average	0.9744	0.92165	0.97398	0.97537
stdDev	0.00093	0.02196	0.00186	0.00012
	t-tests (alpha = 0.05)			
reject null?	yes	yes	yes	baseline
significance	4.59E-05	2.71E-13	0.002	baseline
ci lower	0.0005	0.0438	0.0005	baseline
ci upper	0.0014	0.0637	0.0022	baseline

Table 1: Mutation results with random population initialization.

3. RESULTS AND ANALYSIS

Table 1 presents the results using random population initialization. Though differences in average fitness may appear small when measured in UQI, they appear larger when using alternate error measures, such as mean squared error [7]. Small standard deviations indicate consistent performance across replications. Gaussian mutation achieves the highest average UQI. The other mutation operator results are compared to the Gaussian results using the student's two-sized t-test at 95% confidence ($\alpha = 0.05$). At this confidence level, the Gaussian mutation results are superior at a significant level. The table also provides the significance level (p-value) for each t-test and the 95% confidence level in the difference of means.

While the Gaussian mutation operator provides the best global search, the local and nonuniform operators provide competitive levels of performance. The Cauchy mutation operator performs much worse than the other operators. Due to the large mutations obtained by the Cauchy distribution; the GA does not effectively search promising local areas of the solution space.

Table 2 shows the GA performance under local population

Local Population Initialization Fitness

mutation	NonUniform	Cauchy	Local	Gaussian
	Descriptive Statistics			
average	0.975405	0.973605	0.975365	0.975425
stdDev	0.00008	0.00034	0.00016	0.00006
	t-tests (alpha = 0.05)			
reject null?	no	yes	no	baseline
significance	0.3275	1.87E-24	0.0842	baseline
ci lower	-2.15E-05	0.0017	-0.0091	baseline
ci upper	6.26E-05	0.002	0.1374	baseline

Table 2: Mutation results with local population initialization.

initialization. As before, the standard deviations are small, indicating repeated consistent performance. As with random population initialization, the Gaussian operator provides the best average performance, though the improvement over other operators is no longer clear. T-tests indicate no significant difference in average performance between the Gaussian operator and either the local or nonuniform mutation operators. As before, the Cauchy mutation operator results in the worst performance, but it performs much better when the population is initialized in the neighborhood of the original wavelet coefficients, since the GA is already biased toward a favorable location in the search space. When initializing the population in the local neighborhood of the DWT, the choice of mutation operators is not as important for GA performance. Though results are only presented for the training image, previous research has established that evolved filters continue to outperform the original DWT filters when applied to a diverse set of independent training images [5, 7].

4. CONCLUSIONS

GAs using real-coded representations depend heavily upon the ability of their variation operators (mutation and recombination) to provide a sufficient sampling of very large search spaces. Real-valued operators that do not sufficiently sample the search space often exhibit deteriorated performance [6]. Hence, the selection of appropriate variation operators is critical when designing GAs for military-grade algorithm development. Among the tested mutation operators, Gaussian mutation results in the best global search for evolved image reconstruction transforms. Nonuniform, local, and Gaussian mutation all perform well for searching for improved transforms in GAs initialized near the original DWT reconstruction coefficients. Whether performing a global search or refining local solutions, the Gaussian mutation operator using smaller mutation step probabilities in later generations allows for good search of the solution space using GAs.

The next step in this research will be to evaluate advanced crossover operators for real-valued recombination, such as the unimodal normal distribution crossover (UNDX), the blend crossover (BLX-α), simulated binary crossover (SBX), or simplex crossover (SPX). By sampling a wider area of the search space, such crossover operators outperform traditional crossover operators for real-valued problems in epistatic environments [6] and may result in improved evolution of image and signal transforms.

Military systems involving the acquisition and analysis of large amounts of data require advanced algorithms able to retain critical information at high compression levels when subject to quantization. GAs aid in the development of such algorithms. Though GAs are notoriously difficult to tune and operator selection may be difficult, it is critical to select operators providing fast search for high quality solutions. By replicating small GA tests for evolved reconstruction filters, we are able to identify operators providing strong performance. In the future, we will apply these operators in the evolution of transform filters of longer length in GAs evolved for many generations as we develop military grade signal and image processing algorithms.

Acknowledgment

The authors acknowledge the financial support of the United States Air Force Office of Scientific Research as well as the financial support and technical discussions of Patrick Marshall of the Air Force Research Laboratory, Information Directorate.

5. REFERENCES

[1] I. Daubechies. *Ten Lectures on Wavelets*. SAIM, 1992.

[2] D. B. Fogel and J. W. Atmar. Comparing genetic operators with gaussian mutations in simulated evolutionary processes using linear systems. *Biological Cybernetics*, 63(2):111–114, 1990.

[3] Google Inc. *Google Earth software*, November 2005. Build 3.0.0762.

[4] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, Berlin, 1992.

[5] F. Moore, P. Marshall, and E. Balster. Evolved transforms for image reconstruction. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC'05)*, volume 3, pages 2310–2316. IEEE, 2005.

[6] I. Ono, H. Kita, and S. Kobayashi. A real-coded genetic algorithm using the unimodal normal distribution crossover. In A. Ghosh and S. Tsutsui, editors, *Advances in Evolutionary Computing*, pages 213–237. Springer, 1993.

[7] M. R. Peterson, G. B. Lamont, and F. Moore. Improved evolutionary search for image reconstruction transforms. In *Proc. IEEE Congress on Evolutionary Computation, to appear*, 2006.

[8] B. E. Usevitch. A tutorial on modern lossy wavelet image compression: foundations of jpeg 2000. *IEEE Signal Processing Magazine*, pages 22–35, September 2001.

[9] Z. Wang and A. C. Bovick. A universal image quality index. *IEEE Signal Processing Letters*, 9(3):81–84, March 2002.

[10] X. Yao and Y. Liu. Fast evolutionary programming. In *Proceedings of the* 6^{th} *Annual Conference on Evolutionary Programming*, pages 451–460, 1996.