Revolutionary Image Compression and Reconstruction via Evolutionary Computation, Part 2: Multiresolution Analysis Transforms

FRANK MOORE and BRENDAN BABB
Mathematical Sciences Department
University of Alaska Anchorage
3211 Providence Dr., Anchorage, AK 99508
USA

Abstract: - Previous research demonstrated that a genetic algorithm (GA) can utilize supercomputers to evolve image compression and reconstruction transforms that reduce mean squared error (MSE) by more than 22% (1.126 dB) under conditions subject to quantization, while continuing to average the same amount of compression as the Daubechies-4 (D4) wavelet. This paper describes subsequent research that extends our GA to evolve multi-resolution analysis (MRA) transforms. Test results indicate that our evolved MRA transforms can reduce MSE by an average of more than 10% (0.50 dB) at three levels of decomposition. This result substantially improves upon state-of-the-art MRA transforms for compression and reconstruction applications subject to quantization error.

Key-Words: - wavelets, genetic algorithms, image compression, quantization, multiresolution analysis

1   Introduction

Wavelets [1] are widely used for applications requiring signal compression and reconstruction. Wavelets may be described by four sets of floating-point coefficients: $h1$ (Lo_D) and $g1$ (Hi_D) are the sets of wavelet and scaling numbers for the (forward) discrete wavelet (decomposition) transform (DWT), while $h2$ (Lo_R) and $g2$ (Hi_R) define the wavelet and scaling numbers for the inverse (reconstruction) transform (DWT^{-1}). Fig. 1 lists coefficients for the Daubechies-4 (D4) [1] DWT.

$$h1 = \{-0.1294, 0.2241, 0.8365, 0.4830\}$$
$$g1 = \{-0.4830, 0.8365, -0.2241, -0.1294\}$$
$$h2 = \{0.4830, 0.8365, 0.2241, -0.1294\}$$
$$g2 = \{-0.1294, -0.2241, 0.8365, -0.4830\}$$

Fig. 1. D4 Wavelet Transform Coefficients.

A two-dimensional (2D) DWT of a discrete input image **f** with M rows and N columns is computed by first applying the one-dimensional (1D) subband transform defined by the coefficients from sets $h1$ and $g1$ to the columns of **f**, and then applying the same transform to the rows of the resulting signal. Similarly, a 2D DWT^{-1} is performed by applying the 1D defined by sets $h2$ and $g2$ first to the rows and then to the columns of a previously compressed signal.

A one-level DWT decomposes **f** into M/2-by-N/2 subsignal $\mathbf{a}^1$, $\mathbf{h}^1$, $\mathbf{d}^1$, and $\mathbf{v}^1$, where $\mathbf{a}^1$ is the trend subsignal of **f** and $\mathbf{h}^1$, $\mathbf{d}^1$, and $\mathbf{v}^1$ are its first horizontal, diagonal, and vertical fluctuation subsignals, respectively. Using the MRA scheme [3], a one-level DWT may be repeated $k \leq \log_2(\min(M, N))$ times. The size of the trend subsignal $\mathbf{a}^i$ at level i of decomposition is $1/(4^i)$ times the size of the original image **f**. Nevertheless, the trend subsignal will typically be much larger than any of the fluctuation subsignals; for this reason, the MRA scheme computes a k-level DWT by recursively applying a one-level DWT to the rows and columns of the discrete trend subsignal $\mathbf{a}^{k-1}$. Similarly, a one-level DWT^{-1} may be applied k times to reconstruct an approximation of the original M-by-N image **f**.

The process of mapping the intensity values of a grey-scale image onto a smaller range of possible values is known as quantization. Reducing the numerical precision of each sampled value allows quantized images to be more easily compressed. The corresponding dequantization step, $\mathbf{Q^{-1}}(q)$, produces signal **γ'** that differs from the original signal **γ** according to a distortion measure **ρ**, which may be

computed as the average of the squared error for each sample (i.e., the MSE).

The distortion present in images reconstructed by wavelets increases in proportion to quantization. For many digital signal processing applications, quantization is the only significant source of distortion. For applications requiring high-fidelity imagery, such distortion may be unacceptable.

2 Supercomputers Evolve Improved MRA Transforms: One Set of Coefficients for All Levels

The goal of any image compression and reconstruction system is to simultaneously minimize two parameters:

1. The number of bits needed to represent the compressed, quantized, and encoded image, i.e., the file size (FS).
2. The average distortion observed in reconstructed images, i.e., the MSE.

The purpose of the research described in this paper was to determine whether a GA [2] could utilize the enormous processing power of supercomputers to evolve coefficient sets representing non-wavelet MRA transforms capable of outperforming MRA DWTs under conditions subject to quantization error. The following parameters characterize the GA developed to achieve this goal:

1. The maximum number of generations, G.
2. The size of the evolving population, M.
3. The number of multiresolution levels, MR.
4. The image(s) used to train the GA.

Typical values used during this study were G = 500, M = 2000, and MR = 3. In addition, an Information Entropy (IE) measure was used to provide a fast and accurate estimate of FS during fitness evaluation.

A three-level MRA (forward) transform is illustrated in Fig. 2. In a typical image compression and reconstruction application, a single set of coefficients defining a particular wavelet are used at every level of a MRA transform, allowing the resulting data to be much more easily compressed.

Previous investigations [4] established the overall feasibility of extending the GA-based approach described above to evolve MRA

transforms described by a single set of coefficients. Unfortunately, these studies produced transforms whose MSE reductions averaged only 3.1% SE reduction. Therefore, the purpose of the first task addressed by the new research described in this report was to determine whether Arctic Regional Supercomputer Center (ARSC) supercomputers could be used to evolve a single set of coefficients for use at every level of a MRA transform capable of significantly improving upon previous results.
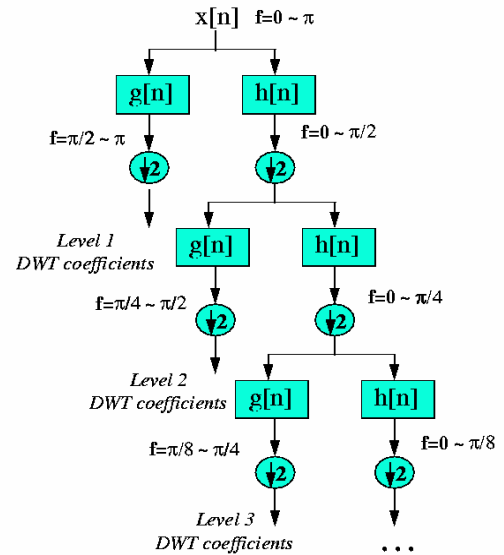

Fig. 2. A Three-level Multiresolution Analysis Transform [8]

Training with the 512-by-512 pixel "zelda.bmp" image and seeding the population with randomly mutated copies of the D4 wavelet, a GA evolved a single set of $g1$, $h1$, $g2$, and $h2$ coefficients that achieved a 10.2% MSE reduction while maintaining an average IE approximately equal to that of the D4. Fig. 3 shows the final coefficients evolved during this run, and lists the percentage difference between each evolved coefficient and the corresponding coefficient from the original D4 wavelet. Note that the greatest percentage change has occurred in the high-frequency coefficients of the reconstruction transform.

These coefficients were used at every level of a three-level MRA transform tested against other 512-by-512 pixel images. The results of these tests (Fig. 4) show an average MSE reduction of over 7.6%. Note that this reduction is more than 2.4 times the

Coefficient Set	Values (Change from D4)
$h1$ (Lo_D)	0.1275, 0.2276, 0.8449, 0.4665 (-1.47%, +1.56%, +1.00%, -3.42%)
$g1$ (Hi_D)	0.4898, 0.8467, -0.2292, -0.1290 (+1.41%, +1.22%, +2.28%, -0.31%)
$h2$ (Lo_R)	.4815, 0.8171, 0.2277, -0.1095 (-0.31%, -2.32%, +1.61%, -15.39%)
$g2$ (Hi_R)	0.1585, -0.1194, 0.7447, -0.3656 (+22.49%, -46.72%, -10.97%, -24.31%)

Fig. 3. Evolved Coefficients and Percentage Change Relative to the D4 Wavelet: One Set of Coefficients Used at Every Level of a Three-level MRA Transform

Image	IE	Improvement (MSE)
airplane.bmp	100.00%	10.49%
baboon.bmp	99.95%	11.55%
barb.bmp	99.95%	14.82%
boat.bmp	100.08%	6.41%
couple.bmp	99.99%	11.66%
fruits.bmp	99.95%	5.79%
goldhill.bmp	100.06%	11.76%
lenna.bmp	99.94%	11.51%
park.bmp	100.03%	9.87%
peppers.bmp	100.08%	13.50%
susie.bmp	99.84%	6.34%
zelda.bmp	100.12%	12.21%

Averages:	100.00%	7.61%

Fig. 4. A Three-level Transform Using a Single Set of Coefficients at Every Level Generalizes Well Against the Test Set of Images.

reduction of the best transform produced prior to utilizing the supercomputer.

To demonstrate the general applicability of the approach, a second run used the "fruits.bmp" image to train a single set of coefficients used at every level of a three-level MRA transform. Test results (Fig. 5) show an average SE reduction of over 7.6% when tested on other 512-by-512 pixel images, while maintaining IE equal to that of the D4 wavelet.

3 Supercomputers Evolve A Different Set of Coefficients for Each of k MRA Transform Levels

Previous work [5] also demonstrated that the GA-based methodology could be used to evolve a different set of coefficients for each level of a MRA transform; for example, a three-level evolved MRA transform derived from the D4 wavelet consists of 48 real-valued coefficients (i.e. 16 coefficients defining the $g1$, $h1$, $g2$, and $h2$ coefficients at each MRA level). The resulting transforms outperformed

both wavelets and evolved transforms described by a single set of coefficients. The average SE reduction for the new transforms (5.7%) was much better than that of evolved transforms described by a single set of coefficients (prior to this research).

The purpose of the second new task addressed by this research was to determine the amount of additional SE reduction that could be achieved by executing our GA-based approach on ARSC supercomputers. Training with the 512-by-512 pixel "zelda.bmp" image and seeding the population with randomly mutated copies of the D4 wavelet, our enhanced GA evolved a three-level MRA transform consisting of 48 real-valued coefficients that achieved a 12.21% SE reduction. Fig. 6 shows the evolved coefficients and the change relative to the D4 wavelet's coefficients. Note that the most significant percentage changes occurred in the high-pass reconstruction transform ($g2$). In addition, significant change occurred in the fourth coefficient of each low-pass reconstruction vector ($h2$). In

Image	IE	Improvement (MSE)
airplane.bmp	100.00%	7.86%
baboon.bmp	99.95%	9.72%
barb.bmp	99.95%	7.23%
boat.bmp	100.08%	7.82%
couple.bmp	99.99%	8.19%
fruits.bmp	99.95%	6.10%
goldhill.bmp	100.06%	8.01%
lenna.bmp	99.94%	7.14%
park.bmp	100.03%	7.40%
peppers.bmp	100.08%	6.56%
susie.bmp	99.84%	7.22%
zelda.bmp	100.12%	8.02%

Averages:	100.00%	7.61%

Fig. 5. A Three-level Transform Using a Single Set of Coefficients at Every Level and Trained on "fruits.bmp" Also Generalize Well Against the Test Set.

contrast, changes to the remaining *h2* coefficients, as well as to the entire *h1* and *g1* coefficient sets of the compression transform, were much smaller.

The evolved coefficients were subsequently tested on several images (Fig. 7). Note that the evolved transform achieved an average SE reduction

of nearly 11% against the test set, while maintaining IE approximately equal to that of the D4 wavelet. This result nearly doubles the average MSE reduction achieved prior to utilizing the supercomputer to run large-scale GA tests.

Set	MRA Level	Values (% Change Relative to D4 Wavelet)
h1 (Lo_D)	1	-0.1278, 0.2274, 0.8456, 0.4664 (-1.24%, +1.47%, +1.09%, -3.44%)
	2	-0.1274, 0.2289, 0.8446, 0.4661 (-1.55%, +2.14%, +0.97%, -3.50%)
	3	-0.1278, 0.2281, 0.8455, 0.4670 (-1.24%, +1.78%, +1.08%, -3.31%)
g1 (Hi_D)	1	0.4791, 0.8474, -0.2347, -0.1278 (-0.81%, +1.30%, +4.73%, -1.24%)
	2	-0.4894, 0.8447, -0.2317, -0.1279 (+1.33%, +0.98%, +3.39%, -1.16%)
	3	-0.4901, 0.8462, -0.2291, -0.1288 (+1.47%, +1.16%, +2.23%, -0.46%)
h2 (Lo_R)	1	0.4811, 0.8152, 0.2274, -0.1069 (-0.39%, -2.55%, +1.47%, -17.39%)
	2	0.4805, 0.8159, 0.2279, -0.1093 (-0.52%, -2.46%, +1.70%, -15.53%)
	3	0.4820, 0.8172, 0.2278, -0.1097 (-0.21%, -2.31%, +1.65%, -15.22%)
g2 (Hi_R)	1	-0.2008, 0.0274, 0.5960, -0.1472 (+55.18%, -87.78%, -28.75%, -69.52%)
	2	-0.1618, -0.1105, 0.6870, -0.3201 (+25.04%, -50.69%, -17.87%, -33.73%)
	3	-0.1572, -0.1495, 0.7861, -0.4033 (+21.48%, -33.29%, -6.03%, -16.50%)

Fig. 6. Different Evolved Coefficients for Each of Three MRA Levels and Percentage Change Relative to the D4 Wavelet.

Image	Original File Size (pixels)	IE	Improvement (MSE)
airplane.bmp	512x512	100.17%	10.49%
boat.bmp	512x512	100.31%	11.55%
boat.bmp	256x256	100.72%	14.82%
baboon.bmp	512x512	100.89%	6.41%
baboon.bmp	256x256	100.70%	13.50%
couple.bmp	512x512	100.43%	11.66%
fruits.bmp	512x512	100.12%	5.79%
goldhill.bmp	512x512	100.34%	11.76%
lenna.bmp	512x512	100.23%	11.51%
park.bmp	512x512	100.47%	9.87%
susie.bmp	512x512	100.25%	6.34%
zelda.bmp	512x512	100.00%	12.21%
-----------	-----------	-----------	-----------
Average Performance:		100.39%	10.49%

Fig. 7. Test Results, Three-level MRA Transform, Different Coefficients at Each Level.

To demonstrate the general applicability of the approach, a second run used the "fruits.bmp" image to train different sets of coefficients for each level of a three-level MRA transform. Fig. 8 lists these coefficients. Test results (Fig. 9) show an average SE reduction of nearly 10.4% in comparison to the D4 when tested on other 512-by-512 pixel images, while maintaining equivalent IE. These results suggest that coefficients trained on representative images generalize well for compression and reconstruction tasks.

Transforms trained on 512-by-512 pixel images also perform very well when tested against smaller images. Fig. 10 tabulates the results of applying coefficients evolved on the 512-by-512 pixel "fruits.bmp" image to a set of 256-by-256 test images. Average MSE reduction exceeded 12.9%, while maintaining an average IE within 0.03% of the D4 wavelet's IE.

6   Conclusions

This paper builds upon previously reported results to clearly establish a new methodology for using GAs to evolve MRA transforms [3] that significantly outperform wavelets under conditions subject to quantization error. Supercomputers have made it possible to achieve much better results than previous investigations. Evolved MRA transforms consistently perform well for images not explicitly anticipated by the training population.

Future investigations will examine the methodology's potential to revolutionize real-world applications currently utilizing wavelets, such as the JPEG2000 image compression standard [10]. In addition, parallel research investigating the use of various crossover and mutation operators on overall system performance ([6], [7]) may be incorporated into the current GA to achieve additional performance improvement. The overall execution time of training runs may be substantially reduced by using representative sub-images for training. Subsignals containing distinctive energy distributions may also be useful in evolving transforms that are capable of highlighting those subsignals when they occur in larger scenes. Techniques for evolving both the number of coefficients in each transform vector, as well as the numerical value of those coefficients, may produce entirely new transforms capable of outperforming any previously defined transforms. An investigation

of alternative evolution-inspired paradigms, such as differential evolution [9], may accelerate the evolutionary process, evolve consistently better transforms, or both.

Coefficients defining the high-pass reconstruction transform (*g2*) consistently exhibited much greater change relative to the corresponding coefficients in the D4 wavelet. Mathematical analysis of these results may provide additional insight into this problem.

References:
[1]   Daubechies, I. 1992. *Ten Lectures on Wavelets*, SIAM.
[2]   Goldberg, D. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.
[3]   Mallat, S. 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation, *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 11(7): 674-693.
[4]   Moore, F. 2005. A Genetic Algorithm for Evolving Multi-resolution Analysis Transforms, *Proceedings of the 5th WSEAS International Conference on Wavelet Analysis and Multirate Systems (WAMUS 2005)*, Sofia, Bulgaria.
[5]   Moore, F. 2005. A Genetic Algorithm for Evolving Improved MRA Transforms, *WSEAS Transactions on Signal Processing*, 1 (1): 97-104, WSEAS Press.
[6]   Peterson, M., G. Lamont, and F. Moore 2006. Evaluating Mutation Operators for Evolved Image Reconstruction Transforms, in *Proceedings of the Eighth Annual Genetic and Evolutionary Computation Conference (GECCO'2006)*, Seattle, WA.
[7]   Peterson, M., G. Lamont, and F. Moore 2006. Improved Evolutionary Search for Image Reconstruction Transforms, in *Proceedings of the 2006 IEEE Congress on Evolutionary Computation (CEC 2006)*, Vancouver, BC, Canada.
[8]   Polikar, R. 1999. The Wavelet Tutorial, Part IV: The Engineer's Ultimate Guide to Wavelet Analysis, `http://users.rowan.edu/~polikar/WAVELETS/WTtutorial.html`
[9]   Price, K., R. Storn, and J. Lampinen 2005. *Differential Evolution: A Practical Approach to Global Optimization*, ISBN: 3-540-20950-6, Springer.
[10] Taubman, D. and M Marcellin 2002. *JPEG2000: Image Compression Fundamentals, Standards, and Practice*, Kluwer Academic Publishers.

Set	MRA Level	Values
h1 (Lo_D)	1	-0.1276, 0.2236, 0.8412, 0.4800
	2	-0.1270, 0.2237, 0.8379, 0.4832
	3	-0.1268, 0.2232, 0.8377, 0.4835
g1 (Hi_D)	1	0.4811, 0.8358, -0.2283, -0.1291
	2	-0.4778, 0.8340, -0.2295, -0.1290
	3	-0.4776, 0.8355, -0.2302, -0.1294
h2 (Lo_R)	1	0.4818, 0.8012, 0.2245, -0.0955
	2	0.4820, 0.8217, 0.2237, -0.1173
	3	0.4820, 0.8219, 0.2243, -0.1171
g2 (Hi_R)	1	-0.1643, 0.0166, 0.6985, -0.2397
	2	-0.1396, -0.1487, 0.8039, -0.4134
	3	-0.1259, -0.1764, 0.8153, -0.4781

Fig. 8. Different Coefficients for Each of Three MRA Levels, Trained on "fruits.bmp".

Image	IE	Improvement (MSE)
airplane.bmp	99.98%	12.62%
baboon.bmp	100.07%	11.86%
barb.bmp	100.04%	2.44%
boat.bmp	100.09%	13.01%
couple.bmp	99.97%	13.13%
fruits.bmp	100.43%	11.66%
goldhill.bmp	100.02%	10.90%
lenna.bmp	99.90%	11.11%
park.bmp	100.00%	11.91%
peppers.bmp	100.02%	7.00%
susie.bmp	99.84%	8.99%
zelda.bmp	100.16%	10.04%
--		
Averages:	100.04%	10.39%

Image	IE	Improvement (MSE)
airplane256.bmp	100.02%	16.10%
baboon256.bmp	100.07%	14.82%
barb256.bmp	100.27%	11.97%
boat256.bmp	99.91%	16.65%
couple256.bmp	99.97%	13.13%
fruits256.bmp	99.93%	0.93%
goldhill256.bmp	99.81%	12.53%
lenna256.bmp	100.12%	15.97%
park256.bmp	100.04%	16.64%
peppers256.bmp	99.99%	12.05%
susie256.bmp	99.97%	12.28%
zelda256.bmp	100.22%	11.93%
--		
Averages:	100.03%	12.92%

Fig. 9. A Three-level Transform Using a Different Coefficients at Each Level and Trained on "fruits.bmp" Also Generalize Well Against the Test Set.

Fig. 10. A Three-level MRA Transforms Using a Different Coefficients at Each Level and Trained on the 512-by-512 Pixel "fruits.bmp" Image Perform Very Well When Tested on 256-by-256 Pixel Images.