Evolved Transforms for Improved Image Compression and Reconstruction under Quantization

FRANK W. MOORE and BRENDAN J. BABB
Mathematical Sciences Department
University of Alaska Anchorage
3211 Providence Dr., Anchorage, AK 99516
USA

Abstract: - Previously reported research efforts demonstrated that a genetic algorithm can evolve coefficients describing transforms that outperform standard wavelets, by reducing the mean squared error (MSE) apparent in reconstructed signals under conditions subject to quantization. This paper describes new results that substantially improve the state-of-the-art in evolved transform performance. Matched forward and inverse transform pairs trained against selected images consistently reduce MSE by more than 22% (1.126 dB) when applied to an arbitrary population of similarly quantized test images, yet still achieve the same amount of compression.

Key-Words: - wavelets, genetic algorithms, image compressions, quantization

1 Introduction

Wavelets [3] are widely used for applications requiring signal compression and reconstruction. For example, the Joint Photographic Experts Group selected the 9/7 Discrete Wavelet Transform (DWT) for Part 1 of the JPEG2000 still image compression standard [12].

An unsigned grey-scale digital image consists of a two-dimensional array of sampled non-negative intensity values $I(x_i, y_j)$, where $0 \le i < X_{max}$ and $0 \le j < Y_{max}$. Quantization is the process of mapping each of these sampled values onto a smaller range of possible values. For example, scalar quantization may be used to map a 16-bit source signal to an 8-bit binary value, resulting in a quantization step of 256:1. Reducing the precision of each sampled value via quantization allows the resulting image to be more easily compressed.

For many digital signal processing applications, quantization is the only significant source of distortion. The corresponding dequantization step, $\mathbf{Q^{-1}}(\boldsymbol{q})$, produces signal $\boldsymbol{\gamma'}$ that differs from the original signal $\boldsymbol{\gamma}$ according to a distortion measure $\boldsymbol{\rho}$. Fig. 1 illustrates the process of compressing, quantizing, encoding, decoding, dequantizing, and reconstructing an image. A variety of techniques may be used to quantify distortion; assuming that quantization errors are uncorrelated, then the aggregate distortion $\boldsymbol{\rho}$ $(\boldsymbol{\gamma}, \boldsymbol{\gamma'})$ in the dequantized signal may be computed as a linear combination of the MSE for each sample.

The distortion present in images reconstructed by wavelets increases in proportion to quantization. Fig. 2 shows an image commonly appearing in the signal processing literature, while Fig. 3 shows the same image after it was compressed, quantized with a quantization step of 64:1, encoded, decoded, dequantized, and reconstructed by a Daubechies-4 (D4) DWT. For medical, scientific, and military applications requiring high-fidelity imagery, such distortion may be unacceptable.

2 Previous Results

A series of projects were reported prior to the research described in this paper. The first of these projects [6] focused entirely upon using a standard genetic algorithm (GA) [4] to evolve a novel set of coefficients for an evolved inverse transform capable of reducing the MSE in a one-dimensional signal previously compressed by the selected wavelet, quantized, encoded, decoded, dequantized, and then reconstructed by an evolved transform having identical structure but different coefficient values. The results were promising, with error reductions exceeding 91% for various sinusoidal signals, and 12% for various ramp signals.

The second project [7] investigated whether the same approach could be successfully applied to images. Test results demonstrated that the GA could evolve inverse transforms capable of reducing MSE by as much as 10.7% in comparison to the selected wavelet.

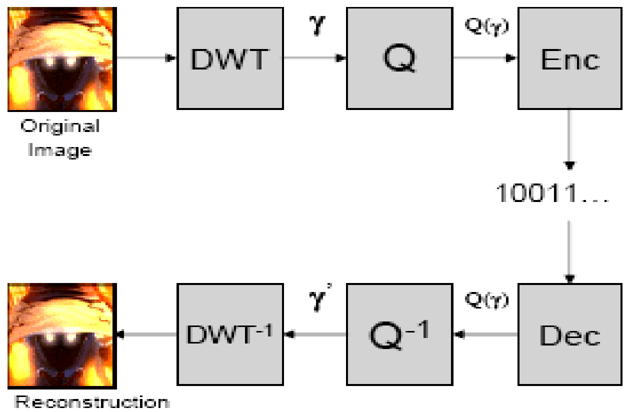

Fig. 1. 1-D Discrete Wavelet Transform Filter with Quantization, Encoding, Decoding, and Dequantization.

Fig.2. Original “zelda.bmp” image.

Fig.3. “Zelda.bmp” image compressed and reconstructed using the D4 wavelet with a quantization step = 64.

Unfortunately, error reductions less than approximately 20% are barely perceptible to the naked eye. The third project [1] therefore focused upon determining whether the simultaneous evolution of coefficients describing matched forward and inverse transform pairs could produce transforms that further reduced MSE in reconstructed images. Initial tests produced transform pairs resulting in up to 90% error reduction compared to the D4 wavelet. Unfortunately, this initial result was meaningless: closer inspection of these results revealed that the forward transform had merely learned to preserve greater accuracy in the reconstructed signal by performing substantially less compression! This error was corrected by integrating file size into the fitness function. Subsequent training runs evolved matched forward and inverse transform pairs capable of more than 23.6% MSE reduction in comparison with the D4 transform, while maintaining a compressed file size less than or equal to the size of the file compressed by the D4 transform.

3 New Results

As these first three projects progressed, it became clear that the amount of computation required to complete even small-scale training runs far outstripped available computational resources. Thus, previous results could only be construed as having established a lower bound on the potential performance gains of evolved transforms in comparison to wavelets. Most of our initial training runs appeared to be making additional evolutionary progress even as they reached a predetermined maximum number of generations. It became clear that only by employing the massive computational resources of supercomputers would it be possible to approach an upper bound on evolved transform performance.

The decision to utilize supercomputer resources required several changes to our approach. The customized wavelet modeling software used for the first three projects was replaced by Matlab’s Wavelet Toolbox, while the hand-crafted GA was replaced by Matlab’s Genetic Algorithm and Direct Search Toolbox. Both tools were integrated into the Matlab Distributed Computing Engine for execution on Arctic Regional Supercomputer Center (ARSC) platforms. In addition, preliminary tests revealed that Information Entropy (IE) provided a consistently accurate prediction of the size of the compressed file; replacing a time-consuming file size calculation algorithm with an IE measure further reduced the computational cost of fitness evaluation.

Fig. 4 tabulates the results of the one supercomputer run, which used the 256-by-256 pixel "zelda.bmp" training image (Fig. 2). These results show a nearly 40% MSE (2.203 dB) reduction for the training image, and an average MSE reduction of nearly 23% (1.126 dB) on test images. In addition, according to the IE measure, compressed FS was less than or equal to the size of the D4 wavelet-compressed FS for every test image.

image	IE % Size	SE %	SE imprv
airplane	95.34	72	28
baboon	94.38	93.2	6.8
barb	97.85	77.12	22.88
boat	98.03	79.28	20.72
couple	96.45	81.61	18.39
fruits	98.06	96.38	3.62
goldhill	98.82	72.91	27.09
lenna	99.11	70.26	29.74
park	97.04	81.64	18.36
peppers	99.61	68.79	31.21
susie	97.57	72.55	27.45
zelda	100	60.22	39.78
	97.68833	77.16333	22.83667

Fig. 4. Transforms trained on "zelda.bmp" significantly outperform the D4 wavelet.

The "zelda.bmp" image reconstructed by the evolved transform is shown in Fig. 5. Figs. 6 and 7 emphasize the amount of error reduction actually achieved by the evolved transform: Fig. 6 shows the difference between the original image and the D4 wavelet-reconstructed image, while Fig. 7 shows the difference between the original image and the evolved transform-reconstructed image. To aid visualization, differences less than 15 were set to zero.

Figs. 8 and 9 present the results of two additional runs of our improved evolutionary system. Transforms trained on a 256-by-256 pixel "fruits.bmp" (Fig. 8) show increased average percentage MSE reduction (1.185 dB) and reduced variance when tested against other images, with equivalent IE. In addition, these transforms appear to generalize across the entire test set more consistently than transforms trained on "zelda.bmp".

Transforms trained on an "airplane.bmp" image of equivalent size exhibit much better error reduction (averaging 2.120 dB) and generalize well across the image test set (Fig. 9); however, the substantially higher levels of IE indicate that the evolved transform could produce larger compressed files than the D4 wavelet. These results corroborate previously reported data [1] indicating the existence of a nearly linear Pareto optimal front [5] describing the tradeoff between file size and MSE in the solution space of evolved transforms.

Fig.5. "Zelda.bmp" image compressed and reconstructed using the evolved transform with a quantization step = 64. Improvements in shading and clarity are obvious to the naked eye.

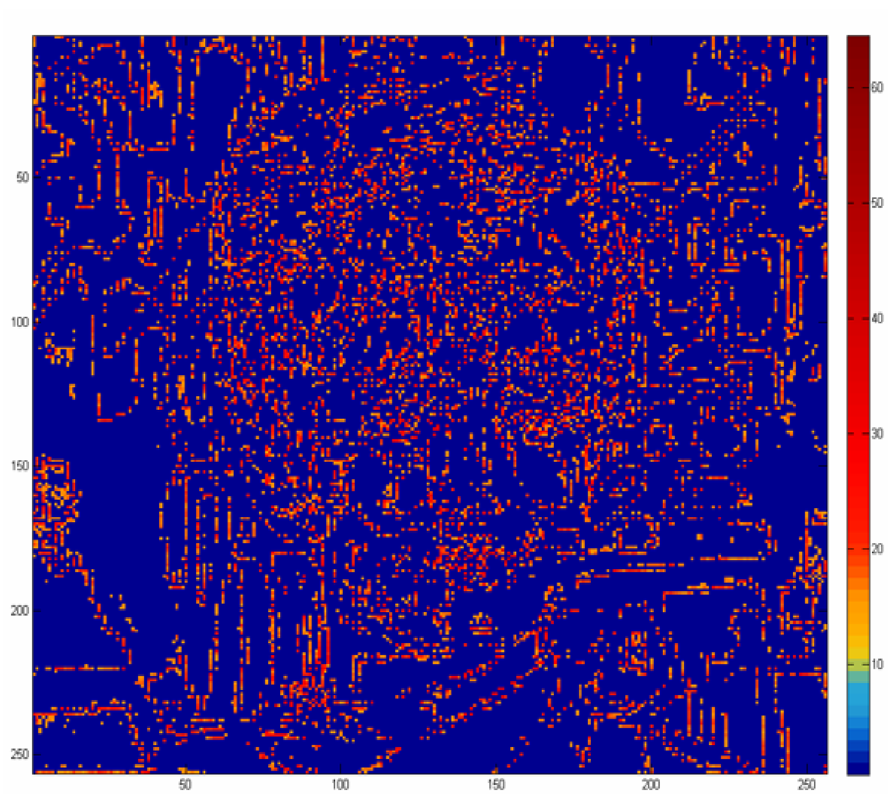

Fig. 6. Differences between the original image and the D4 wavelet-reconstructed image are easily observed.

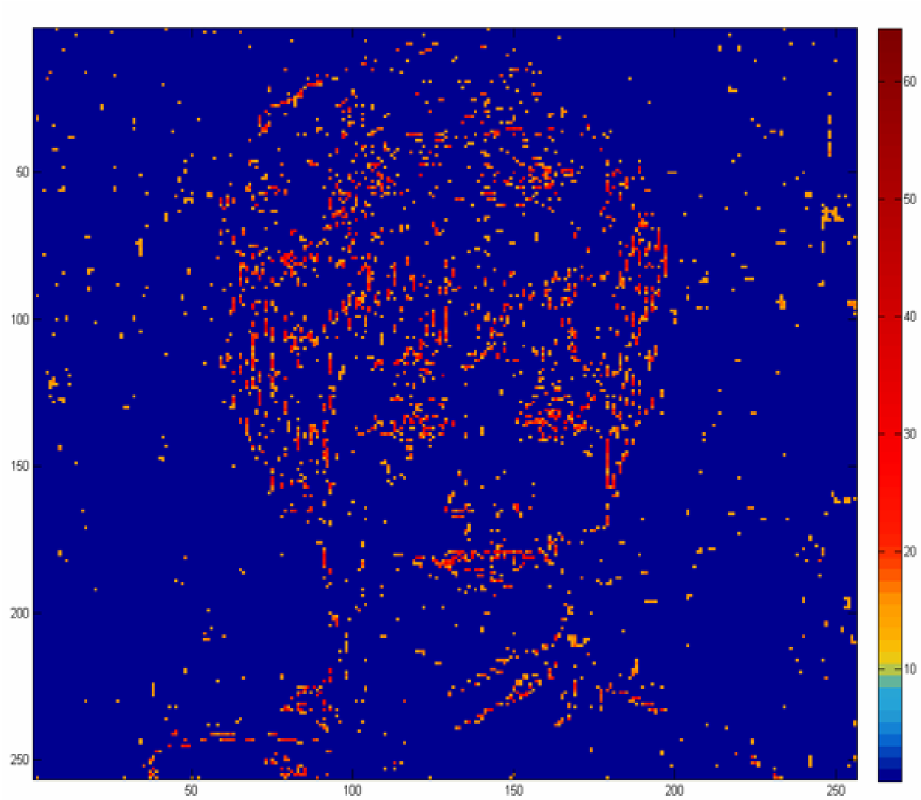

Fig. 7. Differences between the original image and the evolved transform-reconstructed image are much less apparent.

image	IE % Size	SE %	SE imprv
airplane	96.26	72.7	27.3
baboon	98.8	85.07	14.93
barb	100.47	77.72	22.28
boat	99.06	77.34	22.66
couple	100	77.67	22.33
fruits	100	74.82	25.18
goldhill	100.97	73.27	26.73
lenna	100.05	76.75	23.25
park	100.76	86.72	13.28
peppers	101.05	69.02	30.98
susie	100.02	74.45	25.55
zelda	101.51	67.95	32.05
	99.9125	76.12333	23.87667

Fig. 8. Transforms trained on "fruits.bmp" also outperform the D4 wavelet for quantization = 64.

image	IE % Size	SE %	SE imprv
airplane	99.98	57.86	42.14
baboon	105.88	68.6	31.4
barb	105.56	66.09	33.91
boat	105.39	61.73	38.27
couple	105.35	62.55	37.45
fruits	105.24	64.61	35.39
goldhill	105.58	61.93	38.07
lenna	104.47	56.6	43.4
park	104.87	65.17	34.83
peppers	105.72	56.49	43.51
susie	104.12	57.4	42.6
zelda	106.19	57.48	42.52
	104.8625	61.37583	38.62417

Fig. 9. Transforms trained on "airplane.bmp" also outperform the D4 wavelet for quantization = 64.

4 Conclusions

This paper builds upon previously reported results to clearly establish a new methodology for using GAs to evolve transforms that significantly outperform wavelets under conditions subject to quantization error. Our latest results clearly demonstrate that images compressed and reconstructed via evolved transforms exhibit error reductions substantial enough to be clearly visible to the human eye. Furthermore, the evolved transforms consistently perform well for images not contained in the training population.

Future research will utilize the power of supercomputers in an attempt to evolve multi-resolution analysis (MRA) transforms [8] exhibiting similarly large error reduction. An investigation into the methodology's potential to revolutionize real-world applications currently utilizing wavelets, such as the FBI fingerprint compression standard [2] and the JPEG2000 image compression standard [12], is underway. In addition, parallel research investigating the use of various crossover and mutation operators on overall system performance ([9], [10]) 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. Sub-images containing distinctive energy distributions may also be useful in evolving transforms that are capable of highlighting those sub-images 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 reveal the existence of entirely new transforms capable of outperforming any previously defined transforms. Finally, the use of alternative evolution-inspired paradigms, such as differential evolution [11], may accelerate the evolutionary process, evolve consistently better transforms, or both.

References:

[1] Babb, B., S. Becke, and F. Moore 2005. "Evolving Optimized Matched Forward and Inverse Transform Pairs via Genetic Algorithms", in *Proceedings of the IEEE International Midwest Symposium on Circuits and Systems (MWSCAS 2005)*, Cincinnati, OH, 8/07-10, 2005.

[2] Bradley, J., C. Brislawn, and T. Hopper 1993. The FBI Wavelet/Scalar Quantization Standard for Gray-Scale Fingerprint Image Compression, *SPIE Vol. 1961: Visual Information Processing II (1993)*: 293-304, SPIE.

[3] Daubechies, I. 1992. *Ten Lectures on Wavelets*, SIAM.

[4] Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.

[5] Horn, J., N. Nafpliotis, & D. E. Goldberg 1994. A Niched Pareto Genetic Algorithm for Multiobjective Optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence, Volume 1, 1994 (ICEC '94)*. Piscataway, NJ: IEEE Service Center, 82-87.

[6] Moore, F. 2005. "A Genetic Algorithm for Optimized Reconstruction of Quantized Signals", in *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC 2005)*, Edinburgh, Scotland, 9/02-05, 2005.

[7] Moore, F., P. Marshall, and E. Balster 2005. "Evolved Transforms for Image Reconstruction", in *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC 2005)*, Edinburgh, Scotland, 9/02-05, 2005.

[8] Moore, F. 2006. "Evolved Multi-resolution Analysis Transforms for Optimized Image Compression and Reconstruction under Quantization", in *WSEAS Transactions on Computers*, 1(1): 97-104, WSEAS.

[9] Peterson, M., G. Lamont, and F. Moore 2006 (2). "Evaluating Mutation Operators for Evolved Image Reconstruction Transforms", in *Proceedings of the Eighth Annual Genetic and Evolutionary Computation Conference (GECCO'2006)*, 7/08-12, 2006, Seattle, WA.

[10] Peterson, M., G. Lamont, and F. Moore 2006 (1). "Improved Evolutionary Search for Image Reconstruction Transforms", in *Proceedings of the 2006 IEEE Congress on Evolutionary Computation (CEC 2006)*, Vancouver, BC, Canada, 7/16-21, 2006.

[11] Price, K., R. Storn, and J. Lampinen 2005. *Differential Evolution: A Practical Approach to Global Optimization*, ISBN: 3-540-20950-6, Springer.

[12] Taubman, D. and M Marcellin 2002. *JPEG2000: Image Compression Fundamentals, Standards, and Practice*, Kluwer Academic Publishers.