Proceedings of the 2007 IEEE Symposium on Computational
Intelligence in Image and Signal Processing (CIISP 2007)

Evolved Multiresolution Analysis Transforms for Improved Image Compression and Reconstruction under Quantization

Brendan J. Babb[1], Frank W. Moore[1], and Pat Marshall[2]
[1]Mathematical Sciences Department, University of Alaska Anchorage (UAA)
3211 Providence Dr., Anchorage, AK 99516
(907) 786-4819
brendanbabb@gmail.com, affwm@uaa.alaska.edu
[2]AFRL/IFTA Bldg. 620
2241 Avionics Circle, Room 3CY104, WPAFB, OH 45433
(937) 255-6548 x 3586
pat.marshall@wpafb.af.mil

Abstract-**The research described in this paper uses a genetic algorithm (GA) to evolve wavelet and scaling coefficients for transforms that outperform discrete wavelet transforms (DWTs) under conditions subject to quantization. Compression and reconstruction transform pairs evolved against a representative training image reduce mean squared error (MSE) by more than 22% (1.126 dB) when subsequently applied to test images at a single level of decomposition, while evolved three-level multiresolution analysis (MRA) transforms average more than 11% (0.50 dB) MSE reduction when applied to test images in comparison to the Daubechies-4 (D4) wavelet, without increasing the size of the compressed file.**

I. Introduction

Modern digital image compression and reconstruction systems, such as the JPEG2000 standard [12], use wavelets [3]. DWTs may be described by four sets of floating-point coefficients: *h1* (Lo_D) and *g1* (Hi_D) are the 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 these coefficients for the D4 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.

Quantization (the process of representing intensity values using a smaller number of bits) allows images to be more easily compressed. Fig. 2 illustrates the process of compressing, quantizing, encoding, decoding, dequantizing, and reconstructing an image. Quantization is often the most significant source of distortion in digital images. Dequantization step $\mathbf{Q^{-1}}(\boldsymbol{q})$ produces an image **γ'** that differs from the original image **γ** according to a distortion measure **ρ**, which in general may be computed as a linear combination of the MSE for each pixel.

The distortion present in images reconstructed by wavelets increases in proportion to quantization. Fig. 3 shows the "zelda.bmp" image after it was compressed, quantized with a quantization step of 64:1, encoded, decoded, dequantized, and reconstructed by a D4 DWT. For medical, scientific, and military applications requiring high-fidelity imagery, such distortion may be unacceptable.

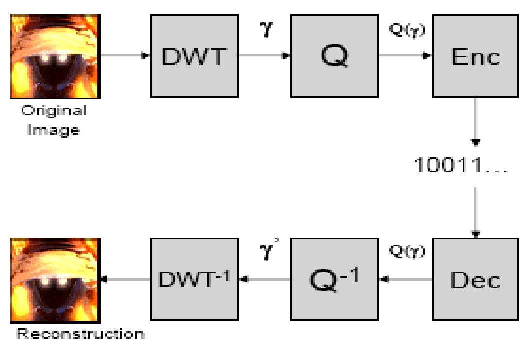

Fig. 2. A Discrete Wavelet Transform Filter with Quantization, Encoding, Decoding, and Dequantization.

II. Previous Results

A series of projects preceded the research described in this paper. The first of these [6] investigated whether a GA could be used to evolve coefficients describing an inverse transform capable of reducing the MSE in reconstructed one-dimensional signals previously compressed by a DWT and subjected to quantization error. Results were promising, with error reductions exceeding 91% for sinusoidal signals. The second project [7] demonstrated that this approach could also be successfully applied to images: a GA evolved inverse transforms capable of reducing MSE by as much as 10.7% in comparison to the selected wavelet. The third project [1] extended this work by simultaneously evolving coefficients describing matched forward and inverse transform pairs. The resulting transforms were capable of more than 20% 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. For each of these projects, the GA seeded the initial population with randomly mutated copies of a selected wavelet; the evolved transforms thus had identical structure to the selected wavelet, but different wavelet and scaling numbers. Note that coefficients are evolved offline (i.e., not in real time), but are subsequently used in real time to compress and reconstruct signals and images not explicitly anticipated by the training population – i.e., evolved coefficients directly replace D4 transform coefficients and are not tailored to specific test images.

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

III. EVOLVED ONE-LEVEL TRANSFORMS

As these previous studies progressed, it became clear that the use of more powerful computer resources would be necessary in order to begin to approach an upper bound on evolved transform performance. The first important goal of this study was to effectively utilize supercomputers to evolve real-valued coefficients describing matched forward and inverse transform pairs capable of outperforming wavelets for image compression and reconstruction tasks subject to quantization error. The customized wavelet modeling software used for the first three projects was abandoned in favor of Matlab’s Wavelet Toolbox, while the hand-crafted GA was reimplemented using 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. 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. 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.

Figs. 5 and 6 emphasize the amount of error reduction actually achieved by the evolved transform: Fig. 5 shows the difference between the original image and the D4 wavelet-reconstructed image, while Fig. 6 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. 7 and 8 present the results of two additional runs of our improved evolutionary system. Transforms trained on a 256-by-256 pixel “fruits.bmp” (Fig. 7) show increased average percentage MSE reduction (1.185 dB) and reduced variance when tested against other images, with equivalent IE. 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. 8); however, 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.

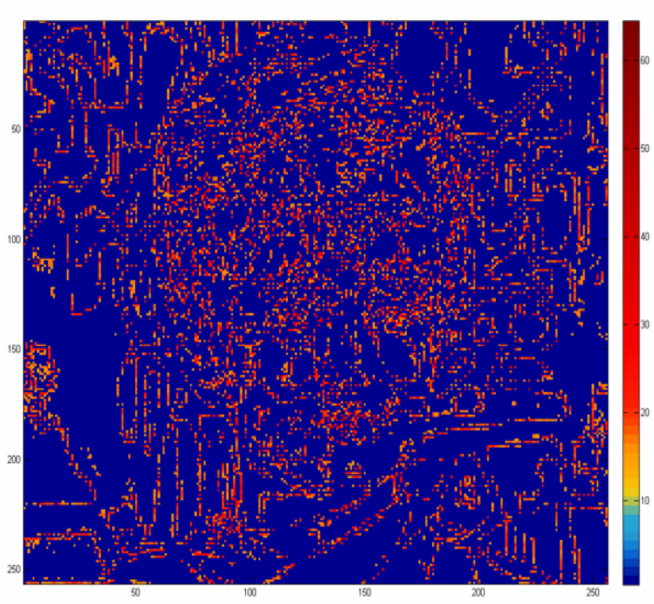

Fig. 5. Differences between the original image and the D4 wavelet-reconstructed image are easily observed.

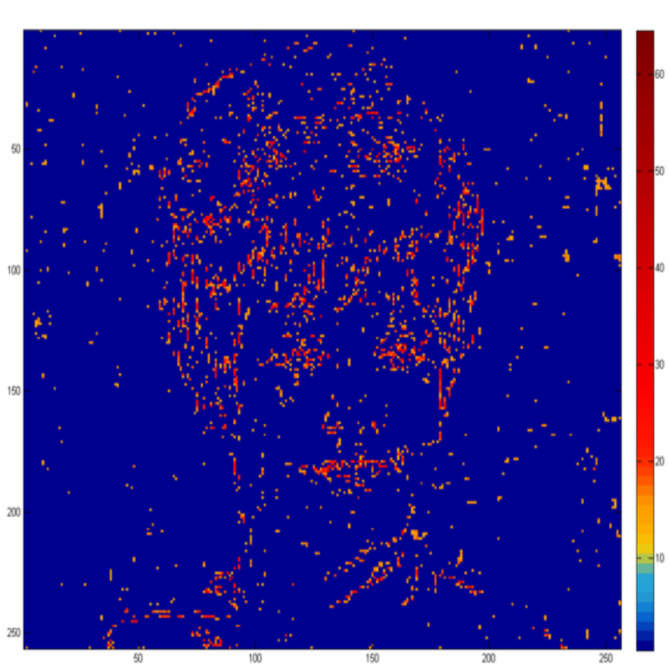

Fig. 6. Differences between the original image and the evolved transform-reconstructed image are much less apparent.

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

IV. Evolved 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 results summarized in the previous section achieve each of these goals. However, in a typical image compression and reconstruction application, a single set of coefficients defining a particular wavelet is used at every level of a MRA transform. Each application of the forward transform achieves additional compression not possible with one-level transforms. For this reason, the next important task of the research described in this paper was to determine whether a GA [4] could 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. IE was again used to provide a fast and accurate estimate of FS during fitness evaluation.

A previous investigation [8] 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% MSE reduction. Therefore, the second important goal addressed by the research described in this paper was to determine whether 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.

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. 9 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. 10) show an average MSE reduction of over 7.6%. Note that this reduction is more than 2.4 times the 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. 11) show an average MSE reduction of over 7.6% when tested on other 512-by-512 pixel images, while maintaining IE equal to that of the D4 wavelet.

V. Evolved MRA Transforms: A Different Set Of Coefficients For Each MRA Level

Previous work [8] 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.

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)	0.4815, 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. 9. Evolved Coefficients and % 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. 10. A Three-level Transform Using a Single Set of Coefficients at Every Level Generalizes Well Against the Test Set of Images.

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. 11. 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.

The third important goal of this research was to determine the amount of additional MSE reduction that could be achieved by using ARSC supercomputers to evolve different sets of coefficients for each level of an MRA transform. Training with the 512-by-512 pixel "zelda.bmp" image and seeding each MRA level of each individual in the population with randomly mutated copies of the D4 wavelet, our enhanced GA evolved a three-level MRA transform that achieved a 12.21% MSE reduction. Fig. 12 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 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.

Evolved coefficients were subsequently tested on several images (Fig. 13). Note that the evolved transform achieved an average MSE 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. 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. Test results (Fig. 14) show an average MSE 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. 15 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.

VI. Conclusions and Future Directions

This paper builds upon previously reported results to establish a new methodology for using GAs to evolve single-level and MRA transforms that significantly outperform wavelets for image compression and reconstruction tasks under conditions subject to quantization error.

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. 12. 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. 13. Test Results, Three-level MRA Transform, Different Coefficients at Each Level.

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%

Fig. 14. A Three-level Transform Using a Different Coefficients at Each Level and Trained on “fruits.bmp” Generalizes Against the Test Set.

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

An investigation into the methodology’s potential to revolutionize real-world applications currently incorporating wavelets, such as the FBI fingerprint compression standard [2] and the JPEG2000 image compression standard [12], is underway. 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. 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.

ACKNOWLEDGMENT

This work was supported in part by a grant of HPC resources from the Arctic Region Supercomputing Center, and is based upon research conducted with support of the US Air Force, the UAA Chancellor’s Challenge Fund, and the UAA Complex Systems Group.

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.