5th WSEAS Int. Conf. on WAVELET ANALYSIS and MULTIRATE SYSTEMS, Sofia, Bulgaria, October 27-29, 2005 (pp12-16)

Evolved Multi-resolution Transforms for Optimized Image Compression and Reconstruction under Quantization

FRANK W. MOORE
Mathematical Sciences Department
University of Alaska Anchorage
CAS 154, 3211 Providence Dr., Anchorage, AK 99508
USA

Abstract: - State-of-the-art image compression and reconstruction techniques utilize wavelets. Recently published research demonstrated that a genetic algorithm (GA) is capable of evolving non-wavelet transforms that consistently outperform wavelets when applied to a broad class of images under conditions subject to quantization error. This paper describes new results that build upon previous research by demonstrating that a GA can evolve a single set of coefficients describing a matched forward and inverse transform pair that can be used at each level of a multi-resolution analysis (MRA) transform to simultaneously minimize the compressed file size (FS) and the squared error (SE) in the reconstructed file. Test results indicate that the benefits of using evolved transforms instead of wavelets increases in proportion to quantization level. Furthermore, coefficients evolved against a single representative training image generalize to effectively reduce SE for a broad class of reconstructed images.

Key-Words: - Wavelets, Genetic Algorithms, Quantization, Multi-resolution Analysis, Image Processing

1 Introduction

Since the late 1980s, engineers, scientists, and mathematicians have used wavelets [2] to solve a wide variety of difficult problems, including fingerprint compression, signal denoising, and medical image processing. Recent adoption of the Joint Photographic Experts Group's JPEG2000 standard [7] has established wavelets as the principal methodology for image compression and reconstruction. Wavelets may be described by four sets of coefficients:

1. *h1* is the set of wavelet numbers for the (forward) discrete wavelet transform (DWT).
2. *g1* is the set of scaling numbers for the DWT.
3. *h2* is the set of wavelet numbers for the inverse DWT (DWT^{-1}).
4. *g2* is the set of scaling numbers for the DWT^{-1}.

For the Daubechies-4 (D4) wavelet, these sets consist of the following floating-point coefficients:

$h1 = \{-0.1294, 0.2241, 0.8365, 0.4829\}$
$g1 = \{-0.4830, 0.8365, -0.2241, -0.1294\}$
$h2 = \{0.4830, 0.8365, 0.2241, -0.1294\}$
$g2 = \{-0.1294, -0.2241, 0.8365, -0.4830\}$

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 ([7], p. 428). Similarly, a 2D DWT^{-1} is performed by applying the 1D DWT^{-1} 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 subimages $\mathbf{h}^1$, $\mathbf{d}^1$, $\mathbf{a}^1$, and $\mathbf{v}^1$, where $\mathbf{a}^1$ is the trend subimage of **f** and $\mathbf{h}^1$, $\mathbf{d}^1$, and $\mathbf{v}^1$ are its first horizontal, diagonal, and vertical fluctuation subimages, respectively. Using the multi-resolution analysis (MRA) scheme [4], a one-level DWT may be repeated $k \leq \log_2(\min(M, N))$ times. The size of the trend signal $\mathbf{a}^i$ at level *i* of decomposition is $1/4^i$ times the size of the original image **f** (e.g., a three-level transform produces a trend subimage $\mathbf{a}^3$ that is $1/64^{th}$ the size of **f**). Nevertheless, the trend subimage will typically be much larger than any of the fluctuation subimages; 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 signal $\mathbf{a}^{k-1}$. Similarly, a one-level DWT^{-1} is applied k times to reconstruct an approximation of the original M-by-N signal **f**.

Quantization is the most common source of distortion in lossy image compression systems. Quantization refers to the process of mapping each of the possible values of a given sampled signal **y** onto a smaller range of values Q(**y**). The resulting reduction in the precision of data allows a quantized signal *q* to be much more easily compressed. The corresponding dequantization step, $Q^{-1}(q)$, produces signal $\hat{\mathbf{y}}$ that differs from the original signal **y** according to a distortion measure ρ. A variety of techniques may be used to quantify distortion; however, if we assume that quantization errors are uncorrelated, then the aggregate distortion $\rho(\mathbf{y}, \hat{\mathbf{y}})$ in the dequantized signal may be computed as a linear combination of SE for each sample.

2 The Genetic Algorithm

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

1. The number of bits needed to represent the compressed image produced by the forward transform (i.e., the *FS*).
2. The distortion observed in the reconstructed image produced by the corresponding inverse transform (i.e., the *SE*).

The purpose of the research described by this paper was to determine whether a GA [3] could be used 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 = 2000.
2. The size of the evolving population M = 2000.
3. The number of multi-resolution levels MR = 3.
4. The probability of crossover p_c = 90%.
5. The probability of mutation p_m for any candidate solution was initialized to a user-specified minimum. If the current generation failed to identify a new best-of-run solution, p_m was increased by a selected increment up to a user-specified maximum mutation rate. The training runs described in this paper used $\min(p_m)$ = 2%, $\max(p_m)$ = 20%, and a 2% increment.
6. The GA trained each transform using a representative 128-by-128-pixel subimage of the standard 512-by-512-pixel "couple.bmp" image. This subimage was chosen based on the results of previous investigations which demonstrated that the resulting evolved transforms generalized well for other images in the test set.

For each of the tests described in this paper, each candidate solution specified the floating-point coefficients for sets *g1*, *h1*, *h2*, and *g2*. The GA seeded the initial population (generation 0) with one exact copy and M-1 randomly mutated copies of the D4 wavelet. Thus, sets *g1*, *h1*, *h2*, and *g2* of every individual in the population each contained precisely four coefficients. After fitness evaluation, the individual with the best fitness value was copied into position 0 of the next generation, while the remaining M-1 positions were populated using tournaments of a user-specified number of randomly selected individuals from the current generation.

Next, the GA performed single-point crossover on adjacent pairs of individuals with probability p_c. The crossover operator randomly selected one of the four coefficient sets, and then randomly selected a crossover point within that set. The coefficients appearing at or below the selected crossover point in the selected coefficient set from each parent were exchanged to create two new candidate solutions.

Finally, mutation was performed on each individual with probability p_m. For this investigation, mutation consisted of multiplying a randomly selected coefficient from a randomly selected set by a factor randomly selected from a Gaussian distribution between 0.0 and 2.0 and centered upon 1.0. Previous studies suggested that an occasional sign change of coefficients could also be beneficial in reducing SE; therefore, with a (typically very small) user-specified probability, the mutation operator used in this study also negated the mutated coefficient.

3 Fitness

This study utilized two key quantities to measure fitness:

1. File Size Ratio (FSR) = FS / (the size of the file compressed by the forward transform)
2. Error Ratio (ER) = SE / (the SE in the image reconstructed by the inverse transform)

Previous research [1] established the existence of a nearly linear Pareto-optimal front describing the trade-off between these two conflicting criteria. For this study, the fitness of each candidate solution against a particular image from the training set was measured as follows:

1. First, the GA used the forward transform coefficients specified by the candidate solution to compress the image.
2. Next, compressed image was quantized using the quantization step defined for the current training run, encoded, decoded, and dequantized.
3. Finally, the GA reconstructed the image using the inverse transform coefficients specified by the candidate solution, and calculated the FSR and ER.

Given a training population consisting of one or more images, this study used the following algorithm to estimate the fitness of a given candidate solution:

```
fitness = 0;
for each training image
   if (FSR > 1.0 && ER > 1.0)
      fitness += FSR^A + ER^B;   // case 1
   else if (FSR > 1.0 && ER ≤ 1.0)
      fitness += FSR^C + ER;     // case 2
   else if (FSR ≤ 1.0 && ER > 1.0)
      fitness += FSR + ER^D;     // case 3
   else
      fitness += FSR^E + ER^F;   // case 4
```

Here, A, B, C, D, E, and F are user-specified constants greater than 1.0. (For this study, A = B = C = D = 8 and E = F = 16.) Lower fitness values are better. Cases 1 and 2 thus explicitly penalize the fitness of transforms that increase the size of the compressed file; cases 1 and 3 penalize transforms that result in higher SE; and case 4 explicitly rewards transforms that simultaneously reduce both compressed file size and SE, relative to the wavelet.

4 One Transform for All MRA Levels

Previous research focused upon evolving coefficients for an inverse non-wavelet transform ([6], [5]) or a matched forward and inverse non-wavelet transform pair [1] that reduced mean SE (MSE) relative to the performance of a standard wavelet transform applied to the same images under conditions subject to a quantization step of 64. The resulting transforms consistently reduced MSE by as much as 25% when applied to images from both the training and test sets. Unfortunately, none of these previous studies involved MRA: instead, coefficients were optimized only for one-level image decomposition and/or reconstruction transforms. Subsequent testing demonstrated that the performance of these transforms degraded substantially when subsequently tested in a multi-resolution environment.

In practice, virtually all wavelet-based compression schemes entail several stages of decomposition. Typical wavelet-based MRA applications compress a given image by recursively applying the *h1* and *g1* coefficients a defining single DWT at each of k levels. Image reconstruction requires k recursive applications of the *h2* and *g2* coefficients defining the corresponding DWT^{-1}. The JPEG2000 standard allows between $0 \leq k \leq 32$ DWT stages; near-optimal performance on full-resolution images is reported for D = 5 levels ([7], p. 429).

The first goal of this research effort was to determine whether a GA could evolve a single set of coefficients for a matched evolved forward and inverse transform pair satisfying each of the following conditions:

1. The evolved coefficients were intended for use at each and every level of decomposition by a matched multi-level transform pair.
2. The evolved forward transform produced compressed files whose size was less than or equal to those produced by the DWT.
3. When applied to the compressed file produced by the matching evolved forward transform, the evolved inverse transform produced reconstructed images whose SE was less than or equal to the SE observed in images reconstructed by the DWT^{-1} from files previously compressed by the DWT.

To achieve this goal, three training runs were performed. These runs differed only according to the specified quantization level. Test results (Fig. 1) confirmed the GA's ability to evolve coefficients for a single transform that exhibited optimized performance when applied to every level of an MRA transform. For Test 3, the GA evolved coefficients that simultaneously reduced SE by almost 6.5% while maintaining a compressed file size smaller than that produced by the D4 wavelet. These results, combined with similar observations from previous studies (e.g., [1]), substantiate the following claims:

1. A GA is capable of evolving matched forward and inverse transform pairs that outperform DWTs at a specified quantization level.
2. The performance improvement of evolved transforms over DWTs increases in proportion to the level of quantization.

Fig. 2 tabulates the coefficients produced by the training runs from Fig. 1, and notes the percentage change in each evolved coefficient from sets *g1*, *h1*, *h2*, and *g2*, relative to the corresponding coefficient

Test	Q	File Size / SE (DWT)	File Size / SE (evolved)	Improvement (SE)
1	16	2162 / 447535.38	2161 / 438555.80	-2.006%
2	32	1229 / 1093462.63	1228 / 1047424.95	-4.210%
3	64	667 / 2527851.95	666 / 2364332.55	-6.469%

Fig. 1. Improvement of Evolved Transforms over Wavelets as a Function of Quantization Level.

Test	Set	Coefficients (Percentage magnitude difference from D4 coefficients)
1	*g1*	-0.4831928406, 0.8365163040, -0.2277694276, -0.1289164106 (+0.05%, unchanged, +1.62%, -0.38%)
	h1	-0.1294917987, 0.2242505778, 0.8398953785, 0.4793849332 (+0.06%, +0.05%, +0.40%, -0.74%)
	h2	0.4830777810, 0.8291240048, 0.2251359248, -0.1227483711 (+0.02%, -0.88%, +0.44%, -5.15%)
	g2	-0.1318678078, -0.1988169414, 0.8344765791, -0.4649087239 (+1.90%, -11.30%, -0.24%, -3.74%)
2	*g1*	-0.4851359202, 0.8394985463, -0.2269758897, -0.1264251009 (+0.45%, +3.57%, +1.26%, -2.31%)
	h1	-0.1300256428, 0.2240904941, 0.8398953785, 0.4798481072 (+0.48%, -0.02%, +0.40%, -0.64%)
	h2	0.4845747470, 0.8203178205, 0.2232898873, -0.1133667585 (+0.33%, -1.94%, -0.38%, -12.40%)
	g2	-0.1312233947, -0.1681967819, 0.8352313868, -0.4547615370 (+1.40%, -24.96%, -0.15%, -5.84%)
3	*g1*	-0.5008454816, 0.8365163040, -0.2158388997, -0.1314604618 (+3.70%, unchanged, -3.71%, +1.58%)
	h1	-0.1285400096, 0.2241438680, 0.8377104749, 0.4827317796 (-0.67%, unchanged, +0.14%, -0.05%)
	h2	0.4896825540, 0.8082258125, 0.2183220074, -0.1034099818 (+1.39%, -3.38%, -2.60%, -20.09%)
	g2	-0.1443190513, -0.1399062106, 0.8240345243, -0.4365732803 (+11.52%, -37.58%, -1.49%, -9.61%)

Fig. 2. Evolved Coefficients and Percentage Change from D4 Coefficients: One Transform for All MRA Levels

from the D4 wavelet. Note that, although coefficient sets *g1*, *h1*, *h2*, or *g2* for every candidate solution were initialized to randomly perturbed copies of the coefficients defining the D4 wavelet, 45 of the 48 coefficients (93.75%) have undergone some change during the evolutionary process. This result corroborates previous test data and underscores the fact that the search space immediately adjacent to the D4 wavelet appears to be rich with non-wavelet transforms that may outperform wavelets under conditions subject to quantization error. Close inspection of these coefficients reveals an interesting phenomenon: in general, the greater the amount of quantization, the greater the difference between evolved coefficients and wavelet coefficients. Also interesting is the fact that none of the evolved coefficients differed in sign from the corresponding wavelet coefficient. Whatever benefits the sign change mutation may have had during previous studies (without MRA) appears to have been eliminated during the evolution of a single set of coefficients for the optimized MRA transforms identified during this study.

5 Generalization Properties of Evolved Transforms

The MRA transform coefficients shown in Fig. 2 were evolved using a single representative subimage extracted from “couple.bmp”. The transform from Test 3 was subsequently tested against several widely used images to determine whether it was capable of achieving similar error reduction for images not used during training. Fig. 3 compares the aggregate SE for the evolved transform from Test 3 to that of the D4 wavelet, when tested under conditions subject to identical quantization error (64). The evolved transform outperformed the D4 wavelet for all but one of the test images. This evidence suggests that transforms trained on a representative subimage are capable of exhibiting optimized performance when tested against a broad class of images having similar visual qualities.

6 Conclusions

This research demonstrated the following key points:

1. A GA could evolve coefficients describing a single matched forward and inverse transform pair that was capable of outperforming a similarly structured standard DWT for a specified MRA level.
2. The advantage of using evolved MRA transforms over DWTs increased in proportion to the specified quantization level.
3. MRA transforms evolved against a representative training image also outperformed DWTs when subsequently tested against arbitrary images.
4. Considerable additional testing will be necessary over a variety of training scenarios to determine whether any discernable pattern in the evolved coefficients emerges.

7 Future Directions

The amount of computation needed to establish an upper bound on the performance enhancement to be gained via evolved transforms far exceeded available resources. The results summarized above should be interpreted as having demonstrated the feasibility of using GAs to evolve optimized MRA transforms. Close inspection of training run results indicate that most GA runs were continuing to make evolutionary progress, even as the number of generations approached G. Thus, larger-scale runs may evolve coefficient sets for MRA transforms that result in considerably greater SE reduction for a given class of images.

References:

[1] Babb, B., S. Becke, and F. Moore 2005. Evolving Optimized Matched Forward and Inverse Transform Pairs via Genetic Algorithms, *Proceedings of the 48th IEEE International Midwest Symposium on Circuits and Systems: Cincinnati, OH, August 7-10, 2005*, IEEE Circuits and Systems Society. In press.

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

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

[4] Mallat, S. 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation, *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 11(7): 674-693.

[5] Moore, F. 2005. A Genetic Algorithm for Optimized Reconstruction of Quantized Signals, *Proceedings, 2005 IEEE Congress on Evolutionary Computation*, IEEE (in press).

[6] Moore, F., P. Marshall, and E. Balster 2005. Evolved Transforms for Image Reconstruction, *Proceedings, 2005 IEEE Congress on Evolutionary Computation*, IEEE (in press).

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

Test Image	SE: Evolved Transform	SE: D4 Wavelet	Change
airplane.bmp	93.92	91.74	+2.4%
baboon.bmp	321.18	331.19	-3.0%
barb.bmp	197.80	202.98	-2.6%
boat.bmp	111.81	116.89	-4.3%
couple.bmp	137.36	143.31	-4.2%
fruits.bmp	80.35	83.00	-3.2%
goldhill.bmp	115.76	120.19	-3.7%
lenna.bmp	102.74	106.86	-3.9%
park.bmp	149.84	155.22	-3.5%
susie.bmp	106.66	110.23	-3.2%
zelda.bmp	51.13	53.81	-5.0%

Fig. 3. Evolved Transforms Exhibit Optimized Performance when Tested Against Other Images.