Evolving Optimized Forward and Reverse Transforms using Genetic Algorithms on a Supercomputer

Brendan J. Babb
Mathematical Sciences Department
University of Alaska, Anchorage
1546 'N' St., Anchorage, AK 99501
(907) 223-2718
bbabb@gci.com

ABSTRACT
The use of digital images is increasing all the time in personal digital photography, medical imaging, and fingerprint image databases. The goal of this research is to improve the image quality of a given compressed digital image while maintaining the same file size of the image. Wavelet based image compression is improved upon by using Genetic Algorithms on a supercomputer to evolve transforms that have better image quality after image compression.

Categories and Subject Descriptors
G.1.2 [**Numerical Analysis**]: Approximation – *Wavelets and Fractals;* I.4.2 [**Computing Methodologies**]: Image Processing and Computer Vision – *Compression (Coding)*; I.2.8 [**Computing Methodologies**]: Artificial Intelligence – *Problem Solving, Control Methods, and Search*; G.1.6 [**Numerical Analysis**]; Optimization - *Global Optimization*

General Terms
Algorithms, Experimentation, Performance

Keywords
Evolved Transforms, Wavelets, Genetic Algorithms, Quantization Error, Image Compression, Image Reconstruction.

1. INTRODUCTION
Digital images are widely used today in digital cameras, medical imaging, and fingerprint image databases, for example. Digital images take up a lot of disk space and often lossy compression is used to store the images in an efficient manner. Lossy compression sacrifices some of the image data in the pursuit of a smaller file size for the image. The goal of my research is to improve the image quality while preserving the original file size of an image.

A common method for compressing images uses wavelet transforms and quantization. Wavelets transforms [1] are similar to Fourier transforms in decomposing a signal or image into the frequency domain. Wavelet-based image transformation occurs when the forward-wavelet transform filters the image data and then the reverse-wavelet transform recovers the exact original image data.

Wavelet transforms are currently used in the JPEG2000 standard, which is a very common file format for digital images. After forward-wavelet transformation, quantization (Q) is used to reduce the data in an image, by mapping the image data values into a much smaller subset. Given image values from 0 to 32,767 data would be quantized by dividing each value, (by say, 128) and rounding down to the nearest integer.

The data values would then range from 0 to 255. The resulting quantized data then undergoes lossless compression. Lossless compression allows the exact original data to be reconstructed from the compressed data before being stored to disk or transmitted.

The image is reconstructed by decompressing the stored or transmitted data, dequantizing it (multiplying each value by 128), and performing the reverse-wavelet transform. The resulting image is very similar to the original image but due to the quantization, it is not exactly the same. Image quality is subjective and based on the perception of the human eye, but there are several metrics that attempt to measure image quality. A common metric for measuring the difference between two images is Squared Error (SE), by which each pixel in the original image is subtracted from the corresponding pixel in the reconstructed image and squared and summed over all the pixels in the image. If the SE is 0, then the reconstructed image is identical to the initial image. A smaller SE indicates a better reconstructed image.

I set out to evolve forward and reverse transforms that would perform better than a wavelet-based transforms in terms of SE under quantization.

2. METHODOLOGY
I chose to concentrate on the Daubechies-4 (D4) discrete wavelet transform (DWT), a standard wavelet designed by Ingrid Daubechies, although the research is applicable for any existing DWT [2, 3, 4]. The D4 is usually represented by the scaling numbers *h1* and the set of wavelet numbers *g1* (see Figure 1) and referred to as the forward transforms. The numbers *h2* and *g2* are mirror reflections of *h1* and *g1* and are referred to as the reverse transforms.

GECCO'06, July 8–12, 2004, Seattle, WA, USA.

$$h1 = \left\{ \frac{\sqrt{2} * (1 + \sqrt{3})}{8}, \frac{\sqrt{2} * (3 + \sqrt{3})}{8}, \frac{\sqrt{2} * (3 - \sqrt{3})}{8}, \frac{\sqrt{2} * (1 - \sqrt{3})}{8} \right\}$$

$$g1(n) = -1^n h[P - n]$$

Figure 1. Daubechies-4 (D4) Wavelet Coefficients.

The D4 is made up of 16 real valued coefficients, so the genetic algorithm (GA) [5] consisted of genomes of 16 real values. The genomes were represented by an array of 16 real values. I used a one-point crossover operator that randomly selected one of the 16 values. The mutation operator selected a random value *X* between 0.8 and 1.2, taking the square root and multiplying one of the 16 values by *X*.

The initial population was seeded with the original D4 coefficients and copies of the D4 coefficients that had been multiplied by a random number between 0.9 and 1.1. The thought was that better solutions lived close to the original wavelet.

Image quality and file size are linked. When the image quality improves the file size increases. However, the aim was to ensure that as image quality increased file size remained the same so both factors were measured in the fitness function.

For a given member of the population, the first 8 values of the genome would be used for the forward transform of an image *X*, which would then be quantized, dequantized and finally the reverse transform consisting of the last 8 real values of the genome would be applied. The resulting image *X'* would be compared with *X* and the SE would be computed.

Wavelet transforms of images tend to result in data in which many of the values are close to zero, quantization maps these values to zero [6]. This allows a smaller amount of data to represent the image. Once the image is quantized, the data undergoes lossless compression and is stored in a file or transmitted. Previously I had used a particular compression method for storing the data. The GA might have been optimizing the compression as well as improving the transform. To isolate the improvement of the transforms, I chose to use Shannon's information entropy (IE) [7] of the data versus a particular compression method.

Shannon's information entropy (sometimes referred to as content) is calculated by taking all the values of the data and counting the number of times a value occurs. The probability *p(z)* of a particular value *z* occurring is determined by dividing the count of *z* by the total number of values in the data. The entropy is then calculated by taking the sum of *–p(z) * log(p(z))* over all *z*, and multiplying by the total number of values. The entropy is considered the theoretical lower limit to which an image can be compressed using a frequency-based lossless compression. The assumption is that an evolved transform that produces data with the same information entropy as the original wavelet can be compressed to the same file size regardless of the lossless compression method chosen.

The fitness was then calculated by considering both the IE and SE. I chose to train the GA on the Fruits image which is a 256 by 256 pixel bitmap image of fruits, and part of the standard library of images used in testing image processing. I tested with a quantization of 64. Since the SE and IE could vary based on the image, quantization and transform, I was interested in the improvement over the original wavelet. The ratios of SE and IE improvement were measured in the fitness function.

SE ratio equals the resultant SE for a transform / (the original SE given the original wavelet and given quantization). For example if the original SE was 200 and the new transform resulted in an SE of 150 the SE ratio would be 0.75.

IE ratio equals the resultant IE for a transform / (the original IE given the original wavelet and given quantization). For example if the original IE was 2000 and the resultant IE was 3000 the IE ratio would be 1.50.

As IE increased, the SE would decrease so a penalty was introduced for transforms that would result in a larger IE than the original. A penalty was also introduced for an SE larger than the original SE. The following block shows the fitness function for the GA:

```
If (SE ratio > 1) and (IE ratio > 1)
   then fitness = (SE ratio)^4 +(IE ratio)^4
else if (SE ratio > 1)
   then fitness = (SE ratio)^4 + IE ratio
else if (IE ratio > 1)
   then fitness = SE ratio + (IE ratio)^4
else
   fitness = (SE ratio)^2 + IE
fitness = fitness *1000
```

The supercomputers at the Arctic Region Supercomputing Center were used to run the GA [8] with a population of 10,000 for 1,000 generations. The supercomputer allowed for runs that had previously taken a week to complete, to run in 8 hours. The supercomputer produced better results [9] then previous runs on desktop PCs.

3. RESULTS

I achieved 23.8% (Table 1) average SE ratio improvement (SERI) over the standard D4 wavelet with an average IE less than 100% over all the test images. The GA was trained using the Fruits image. The evolved transforms gave a 25.18% SERI for Fruits. The evolved transforms generalize well. The Airplane image had a 27.30% SERI, larger than Fruits, with the IE 96.26% of the original IE.

Since the IE for airplane was less than 100%, I decided to see if the SE could be improved further if I trained on Airplane. I started from the evolved Fruits transforms and trained on Airplane for 150 generations with a population of 1000. SERI for Airplane was 42.14% with an IE of 99.98%. The average SE ratio improved to 38.62% but the average IE ratio also increased to 104.86%.

Figure 2. This image has undergone the initial D4 wavelet transformation using a quantization of 64 to compress an original photograph of Zelda.

Figure 3. This improved image has undergone compression using the evolved transform with a quantization of 64 resulting in 39.78% Squared Error Ratio Improvement.

Table 1. Improvement from evolved Fruits transform compared to original D4 transform with quantization 64.

Image	IE %	SE %	SE imp. %
Airplane	96.26	72.70	27.30
Baboon	98.80	85.07	14.93
Barb	100.47	77.72	22.28
Boat	99.06	77.34	22.66
Couple	100.00	77.67	22.33
Fruits	100.00	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	69.04	30.98
Susie	100.02	74.45	25.55
Zelda	101.51	67.95	32.05
Averages	**99.91**	**76.12**	**23.88**

Tests on the Zelda image led to the highest SERI at 32.05% but with a 101.51% IE ratio, so I attempted to continue training on Zelda from the evolved Fruit transforms. I used a population of 1000 and evolved for 150 generations, and the resulting SERI for Zelda was 39.78% (see Figures 2 and 3) with an IE ratio of 100%. The average IE ratio was 97.69% and the average SERI was 22.84%.

4. FUTURE RESEARCH

Both the Zelda and Airplane image transforms show individual SERI of close to 40%. Can additional GA runs produce a generalized transform that achieves an average SERI of 40% or greater?

This GA method was evolved over a test bed of photographic images and showed significant image quality improvement. Additional gains might be made using a test bed of fingerprint images or medical images.

The population was seeded with individuals similar to the original wavelet; additional research will investigate starting with random valued populations.

Differential Evolution is a variant of the Genetic Algorithm that has shown quicker convergence on real valued problems and might increase the SERI.

5. ACKNOWLEDGMENTS

This research was supported by a Chancellor's Challenge Fund grant at the University of Alaska Anchorage. Special thanks to Dr. Frank Moore for selecting me to work on this project. Special thanks to Michael R. Peterson at AFIT/ENG for his advice and allowing me to extend his MATLAB code. Supercomputer access and technical assistance were provided by the knowledgeable staff at the Arctic Regional Supercomputer Center in Fairbanks, AK. This work was supported in part by a grant of HPC resources from the Arctic Region Supercomputing Center at the University of Alaska Fairbanks as part of the Department of Defense High Performance Computing Modernization Program.

6. REFERENCES

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

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

[3] Moore, F., P. Marshall, and E. Balster 2004. Adaptive Filtering in the Wavelet Transform Domain via Genetic Algorithms. Final Report, USAF Office of Scientific Research. AFRL-IF-WP-TR-2004-1553.

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

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

[6] Walker, J. S. 1999. *A Primer on Wavelets and their Scientific Applications*, CRC Press.

[7] Shannon, C. E. A Mathematical Theory of Communication. *Bell System Technical Journal*, (July 1948), 623

[8] M. Peterson, G. Lamont, and F. Moore, 2006. Improved Evolutionary Search for Image Reconstruction Transforms. *Proceedings of the 2006 Congress on Evolutionary Computation*, IEEE (in press).

[9] Babb, B., Becke, S. and Moore, F. 2005. Evolving Optimized Matched Forward and Inverse Transform Pairs via Genetic Algorithms. *MWSCAS 2005, The 48th IEEE International Midwest Symposium on Circuits and System.*