Targeted Filter Evolution for Improved Image Reconstruction Resolution

Michael R. Peterson
Dept. of Computer Science and Engineering
Wright State University
Dayton, OH, USA
peterson.7@wright.edu

Gary B. Lamont
Dept. of Electrical and Computer Engineering
U. S. Air Force Institute of Technology
WPAFB, OH, USA
gary.lamont@afit.edu

Frank Moore
Dept. of Mathematical Sciences
University of Alaska Anchorage
Anchorage, AK, USA
affwm@uaa.alaska.edu

Patrick Marshall
AFRL/IFTA
WPAFB, OH, USA
patrick.marshall@wpafb.af.mil

ABSTRACT

Government, commercial, scientific, and defense applications in image processing often require transmission of large amounts of data across bandwidth-limited channels. Applications require robust transforms simultaneously minimizing bandwidth requirements and image resolution loss. Image processing algorithms take advantage of quantization to provide substantial lossy compression ratios at the expense of resolution. Recent research demonstrates that genetic algorithms evolve filters outperforming standard discrete wavelet transforms in conditions subject to high quantization error. While evolved filters improve overall image quality, wavelet filters typically provide a superior high frequency response, demonstrating improved reconstruction near the edges of objects within an image. This paper presents an algorithm to generate transform filters that optimize edge reconstruction, improving object edge resolution by up to 24%. Such filters provide an increased object resolution over standard wavelets and traditionally evolved filters for varied applications of image processing.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods and Search—*Heuristic Methods*;
I.4.5 [**Image Processing and Computer Vision**]: Reconstruction—*Transform Methods*

General Terms

Algorithms, Design

GECCO '07,July 7-11, 2007, London, England, United Kingdom

Keywords

Genetic algorithms, wavelets, image processing

1. INTRODUCTION

Image and signal processing are active areas of defense and scientific research. Satellites and unmanned aerial vehicles (UAVs) potentially collect and transmit huge volumes of data during reconnaissance missions. Sonar and radar systems process huge amounts of sensor data in real time. Deep-space probes require robust data encoding algorithms to compensate for noise induced by electromagnetic interference or solar radiation. In each case, the need to minimize mission cost while maximizing performance motivates the development of compression techniques that simultaneously minimize bandwidth and storage requirements while maintaining maximum signal information.

With these requirements in mind, quantization is often necessary for military and scientific digitial signal processing (DSP) applications. Shannon's theorem places limits on the amount of compression achievable by any lossless encoding algorithm [18]. In order to achieve higher rates of compression through higher energy compaction than lossless encoders allow, algorithms must permit some loss of information. *Quantization* minimizes storage requirements by mapping all values in signal x to a small alphabet of values $Q(x)$ [4]. Smaller alphabets provide greater compression but result in greater data loss. Perfect reconstruction of x from $Q(x)$ is impossible due to the loss of low-order bits [21].

Wavelets [2] are a standard methodology for signal compression algorithms. The discrete wavelet transform (DWT) redistributes the energy in a signal by transforming a time signal into a time-frequency domain. A signal may be compressed by first applying the DWT, followed by quantization, and then by applying entropy coding. Signals are reconstructed in a reverse manner. Most information loss occurs during quantization. Wavelets have become a popular technique for image coding [3] and provide the algorithmic basis for the JPEG 2000 image compression standard [21].

Wavelet performance degrades at high quantization levels; evolved filters replacing wavelet coefficients may improve image resolution [16]. This paper will demonstrate that while filters evolved using previously published techniques improve the overall mean-squared error (MSE) of reconstructed images, the original wavelet filter provides a better response near the edges of objects within images. For reconnaissance and intelligence-gathering applications, the resolution of edges is crucial for identification and analysis. We propose an algorithm for the evolution of filters designed specifically for reconstructing the portions of images near object edges. Combined with traditionally evolved filters that reconstruct the remainder of the image, the resulting filter combinations provide improved object resolution and provide greater MSE reduction than either wavelets or traditionally evolved filters alone provide.

2. BACKGROUND

2.1 The Discrete Wavelet Transform

In ideal conditions, the DWT^{-1} algorithm provides near-perfect reconstruction of a DWT-decomposed image signal. However, for some applications, higher compression ratios are required than can be achieved by DWT decomposition and entropy encoding alone. In such cases, a quantizer transforms the DWT-decomposed signal to a smaller alphabet before entropy encoding. A corresponding dequantizer translates the decoded signal back to the original alphabet before image reconstruction via the DWT^{-1} algorithm. This process is demonstrated in figure 1.

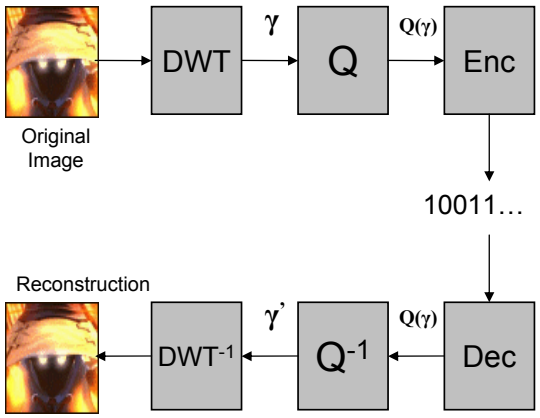

Figure 1: A wavelet-based image compression model

A DWT convolves a signal against specific wavelet instances at various time scales and positions, resulting in a compressed representation of the original signal. The compression is reversed via the corresponding DWT^{-1} by convolving the compressed signal against an inverted order of the original wavelet instances to produce an approximation of the original signal. Wavelets conserve energy and redistribute the bulk of that energy to the "first trend" subsignal [22]. Most of the transformed signal's remaining values outside of the first trend are insignificant and may be eliminated without significant loss of information, providing a favorable compression rate at the expense of perfect reconstruction.

A scaling function $\phi(t)$ and a wavelet function $\varphi(t)$ characterize any DWT as follows:

$$\phi(t) = \sum_n h_n \phi(2t - n) \tag{1}$$

$$\varphi(t) = \sum_n g_n \phi(2t - n) \tag{2}$$

where h_n is the scaling filter's impulse response, g_n is the wavelet filter's impulse response, and n is the translation parameter. The scaling number set $h_1 = \{h_n\}$ and wavelet number set $g_1 = \{g_n\}$ provide coefficients corresponding to the projection of the basis functions for the DWT's low-pass and high-pass filters [2, pp. 54–56]. h_1 and g_1 for the Daub4 wavelet are shown in equations 3 and 4, where $P = 4$ is the size of each set and $n = 1..P$ is the position of the coefficient within the ordered set.

$$h_1 = \{\frac{\sqrt{2}*(1+\sqrt{3})}{8}, \frac{\sqrt{2}*(3+\sqrt{3})}{8}, \frac{\sqrt{2}*(3-\sqrt{3})}{8}, \frac{\sqrt{2}*(1-\sqrt{3})}{8}\} \tag{3}$$

$$g_1 = -1^n * h_1\{P - n\} \tag{4}$$

Sets h_2 and g_2, consisting of mirror images of sets h_1 and g_1, define the DWT^{-1}.The inverse transform reconstructs an approximation of the original signal by convolving the compressed signal using h_2 and g_2.

Let image $\mathbf{f}$ have dimensions M and N. A single level DWT applied to $\mathbf{f}$ results in subimages $\mathbf{a}^1$, $\mathbf{h}^1$, $\mathbf{d}^1$, and $\mathbf{v}^1$, each of size M/2 by N/2. $\mathbf{a}^1$ is the first trend subimage of $\mathbf{f}$, concentrating most of the information in $\mathbf{f}$. The remaining subimages $\mathbf{h}^1$, $\mathbf{d}^1$, and $\mathbf{v}^1$ are its first horizontal, diagonal, and vertical fluctuation subimages. Containing most of the information in $\mathbf{a}^1$ results in improved compression obtained during entropy coding. The small values in the fluctuation subimages require fewer bits to encode. By employing multiresolution analysis (MRA), the DWT may be applied to up to $k \leq \log_2(\min(M, N))$ times. The DWT is recursively applied to $\mathbf{a}^1_{i-1}$, where i is the current level of decomposition. When $i = 1$, $\mathbf{a}^i_0$ is the original image $\mathbf{f}$. By applying multiple decomposition steps, the majority of the energy within the image signal is restricted to smaller trend subimages, minimizing the data to encode. MRA reconstruction occurs in reverse order of decomposition by combining the subimages at level i, applying the DWT^{-1} to obtain $\mathbf{a}^{i-1}_0$, and repeating until the image is reconstructed.

In a typical wavelet coding system, the signal γ emerging from the DWT is quantized (see figure 1). γ is mapped onto a restricted alphabet $Q(\gamma)$. The quantized signal is not changed by entropy encoding and decoding, assuming the use of a lossless encoder, such as standard Huffman encoding.Applying dequantization to $Q(\gamma)$ results in a dequantized signal $\gamma' \neq \gamma$, rendering prefect reconstruction by the DWT$^-$1 impossible [21].

2.2 EC and Image Transforms

In recent years, evolutionary algorithms have been used in conjunction with wavelets for a variety of signal processing applications, including signal approximation [12], signal classification [10], and signal compression and reconstruction [6], [14].

EAs have been used with wavelets for a number of image processing applications. Bruckmann et al. employ a binary

Figure 2: USAF museum satellite image.

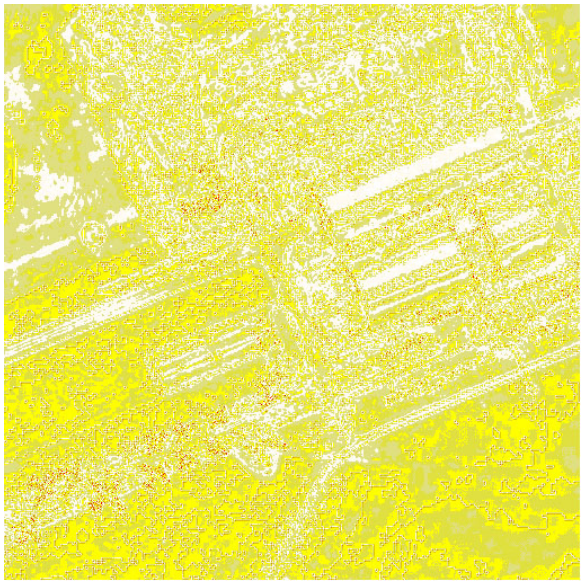

Figure 3: USAF museum db4 wavelet reconstruction error.

GA to evolve subband structures for wavelet packet based image compression [1]. In [17], a GA configures a Kohonen self organizing map (SOM) that, taken with a wavelet-based filter, provides robust image texture classification. Hill et al. evolve a windowed trigonometric function for use in a continuous wavelet transform [8]. In [7], a GA applies the lifting technique [20] to design complementary wavelet filters. The evolved wavelets outperform the standard FBI wavelet [9] for fingerprint image compression.

[13] presents a GA evolves digital filters exploiting MRA by initializing the GA population with values near the original DWT filter and then searching for improved filters in the neighborhood of the original wavelet through a local mutation mechanism. The GA successfully improves image reconstruction both when evolving a single filter for all MRA levels or when evolving unique filters for each level of MRA wavelet decomposition. In related work, a GA evolves only the reconstruction coefficients of a wavelet-based filter to improve image reconstruction in the presence of quantization error [16]. By focusing on the evolution of optimized reconstruction coefficients, the underlying compression rate of the forward transform is unaffected. The resulting filters described in [15, 16] no longer conform to the mathematical properties of wavelets, such as biorthogonality of the filters. Evolved with one or more training images, the resulting filters provide improved reconstruction when applied to images not explicitly represented by the training image population.

The evolutionary approaches of [15, 16] evolve filters designed to reduce the reconstruction error of an entire image. These techniques successfully provide significant MSE reduction to their wavelet counterparts, typically at the expense of increased error in localized portions of images. Consider the satellite image of the U.S. Air Force museum in Dayton, Ohio in figure 2. Figure 3 shows the intensity in error after transformation with a Daubechies-4 (db4) DWT at one MRA level and a quantization step size of 64. Darker pixels indicate greater error, defined as the absolute difference in pixel values between the original image and the reconstructed image. Note that the error is fairly evenly distributed across the entire image. The db4 wavelet achieves 138.13 MSE. Using the GA described in [16] a filter evolved to improve on the db4 reconstruction reduces error in the low-spatial frequency areas of the image at the cost of increased error in the high-spatial frequency areas (i.e. the edges of objects, see figure 4). This image depicts the areas of error between the reconstructed image using the evolved filter. Increased error occurs near the edges of objects, such as the edges of planes and hangars. The evolved filter reduces the MSE to 107.54, an improvement of 22.15%, at the expense of increased error at the edges of objects. Ideally, evolved filters increase the amount of intelligence that can be gathered from reconnaissance images, so this behavior is not desired. This paper presents an approach to reduce reconstruction error near the edges of objects within the reconstructed image without increasing error in the low-spatial frequency areas of the object. Improved reconstruction near edges provides improved object resolution, thus preserving the intelligence that may be gathered from images subject to data loss from high quantization.

3. METHODOLOGY

3.1 Satellite Image Set

Images obtained during reconnaissance missions typically require analysis by intelligence experts. Objects within images must be identified with a high degree of confidence, requiring the highest resolution possible. Mission conditions may result in a loss of resolution due to bandwidth restrictions requiring quantization or data loss due to noise induced by interference. The development of image transforms preserving object resolution to the best possible degree requires training images that may be obtained by satellites or UAVs. Figure 5 shows four images used for GA training in our experiments. Each image is cropped from a high definition

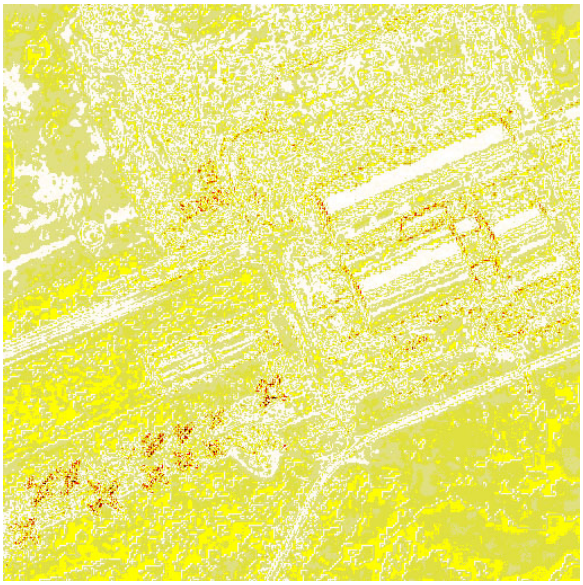

Figure 4: USAF museum evolved filter reconstruction error.

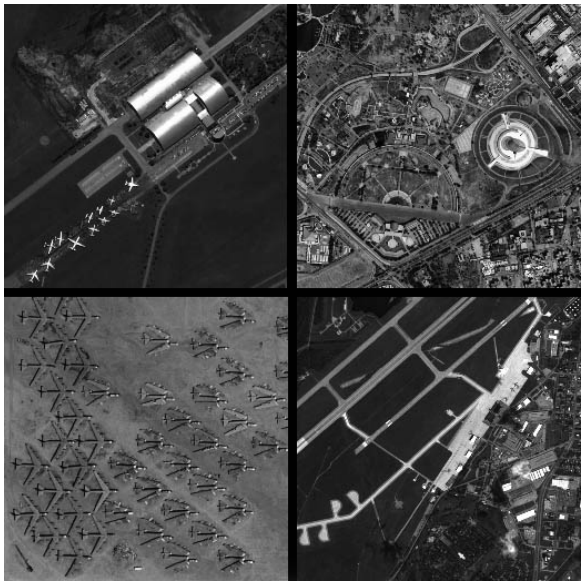

Figure 5: Satellite images used in experiments. Top left: USAF museum. Top right: Baghdad, Iraq. Bottom left: B-52s at Davis-Monthan U.S. Air Force base. Bottom right: Wright-Patterson U.S. Air Force base.

image publicly available using the Google Earth database [5]. The top left image of the U.S. Air Force Museum in Dayton, Ohio contains aircraft, two hangers, and a third under construction. The top right image shows buildings near downtown Baghdad, Iraq. At the bottom right, various B-52 aircraft are seen at the Davis-Monthan U.S. Air Force base near Tucson, Arizona. The final image depicts various buildings and a portion of a runway at Wright-Patterson U.S. Air Force base in Dayton, Ohio. Each is a 512 by 512 pixel black and white image that has been adjusted for maximum contrast. These images are selected because they contain objects whose shapes must be preserved for identification by intelligence experts, and hence provide appropriate training images for the development of transforms preserving object edge resolution.

3.2 Filter Evolution

To demonstrate improvement over previously published techniques, this research employs the GA described in [16], using Matlab's Genetic Algorithm and Direct Search Toolbox. Each experiment evolves a population of 50 filters for 500 generations without early termination. DB4 wavelet reconstruction filters are defined by eight real-valued coefficients. The GA employs a chromosome of eight double precision coefficients, replacing the DB4 DWT^{-1} coefficients to define a new image reconstruction filter. Each generation, the GA copies the two fittest individuals into the next generation. Recombination and mutation are used to create 70% and 30% of the remaining offspring in each generation, respectively. The relatively high rate of mutation is empirically determined from initial experiments.

Recombination consists of Wright's heuristic crossover [23] in which a child lies on the line between the two parents, closer to the parent with the better fitness. Parents are chosen using stochastic uniform selection. This operator is specifically intended for use with real-valued chromosomes. The standard initialization operator randomly creates genes using a random uniform distribution in the range [-1,1]. Mutation adds a random value taken from a Gaussian distribution centered at a randomly chosen parent with a variance of 0.5 at the first generation. The mutation shrinks in successive generations. At generation k, the variance is:

$$var_k = var_{k-1}(1 - .75 * \frac{k}{Gens}) \quad (5)$$

where $Gens$ is the maximum generation. Initially, large variance permits fast exploration of the search space. As the variance shrinks, and the mutation makes very small refinements with increasing probability. The initial population includes one chromosome consisting of original Daub4 reconstruction coefficients. The remaining individuals are copies of the original wavelet coefficients multiplied by a small random factor. Additionally, 5% of the Daub4 coefficients are negated.

The GA only evolves reconstruction filters, always using the original db4 wavelet decomposition filter. Experiments employ a quantization level $q = 64$, meaning that each value in the wavelet-decomposed signal γ is integer divided by 64 with the remainder discarded. Resulting values are dequantized via multiplication by 64 before reconstruction. While the GA identifies improved filters at multiple levels of resolution ([16, 13]), this paper employs a single level of MRA image decomposition and reconstruction to enable faster GA performance.

During fitness evaluation previously decomposed images are reconstructed with a candidate filter using a function for 2-dimensional reconstruction found in Matlab's Wavelet Toolbox. The fitness function determines the similarity of a reconstructed image to the original via mean squared error (MSE) Let $x = \{x_i | i = 1, 2, ..., N\}$ and $y = \{y_i | i = 1, 2, ..., N\}$ represent original and reconstructed images. The

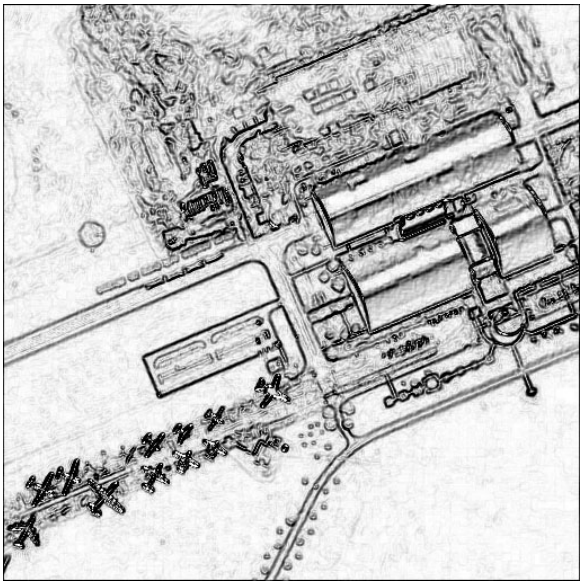

Figure 6: Edge detection algorithm applied to figure 1.

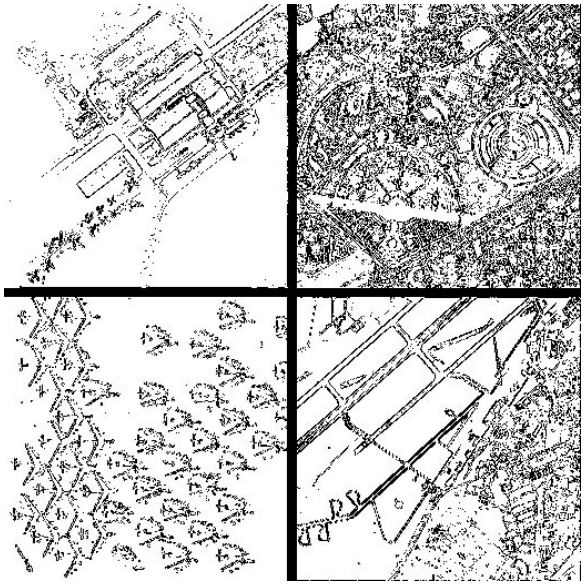

Figure 7: Masks for each test image at threshold 88.

MSE between x and y is:

$$MSE(x,y) = \frac{1}{n}\sum_{i=1}^{n}(x_i - y_i)^2 \quad (6)$$

The GA seeks to improve image reconstruction by minimizing MSE.

3.3 Edge Detection and Mask Generation

Error in images reconstructed with evolved filters occurs near object edges. The first step to counter this effect is to isolate these edges using an edge detection algorithm to identify areas containing significant transitions in pixel intensities. The classic Sobel edge detector performs a 2-D spacial gradient convolution using a pair of 3x3 convolution kernels responding edges running vertically and horizontally relative to the pixel grid [19]. We employ an edge detector developed using a neuro-fuzzy training of a Sugeno-type fuzzy inference system to improve the response in non-pixel axis directions of a Sobel-inspired operator [11]. Figure 6 shows the edges in figure 2 identified by this edge detector.

Once the edges of an image have been isolated, the GA evolves a filter to reconstruct the portions of an image near edges. A binary mask image is created from the edge image by setting a pixel threshold. Pixels darker than the given threshold in the edge image are set to black in the mask, enveloping the edges. The remaining pixels are white. The black portions of the mask are used to select a portion of the original training image to consider during fitness evaluation. Figure 7 shows the masks generated with a threshold of 88 for the four training images. These masks are employed as described in the following section to improve the resolution of reconstructed images near object edges.

3.4 Image Reconstruction

During evolution, the entire image is reconstructed, but fitness is only calculated at the pixel positions located within the black portions of the mask enclosing the edges identified by the edge detection algorithm. This approach forces the GA to evolve a filter that improves image reconstruction near object edges. An image is reconstructed from the db4 decomposed, quantized, and dequantized signal by applying an inverse transform to the entire image using both the filter evolved using the edge-portions of the training image as well as by a filter evolved using the entire image (i.e. not focused on edges). The final image a combination of the two resulting images. Areas of the image near edges indicated by the black portions of the mask are selected from the edge-filter reconstructed image. The remaining portions of the image are selected from the image reconstructed using the traditionally (globally) evolved filter, trained using an entire image. MSE can be lower for the combined image than for the image reconstructed with only the globally evolved filter because error will be reduced near edges. For clarity we will refer to filters evolved using the entire image as globally evolved filters and to filters evolved using the edge-enclosing masks as locally evolved filters.

4. EXPERIMENTS AND DISCUSSION

4.1 Mask Threshold Determination

The creation of the binary mask used to isolate the edge portions of the training image requires a set threshold. This threshold dictates the required strength of the edge detection output for a given pixel position to be considered part of an edge for the mask. In the range $[0-255]$, lower thresholds select fewer areas of the image as edges. Higher thresholds enclose a higher portion of the training image. To determine an appropriate setting for the edge threshold, several GA runs are conducted using the AF museum image from figure one. At a quantization level of 64 and one level of decomposition, the MSE of the image reconstructed using the db4 DWT^{-1} is 138.127. A globally evolved filter, used as a baseline for comparison, achieves a reconstruction MSE

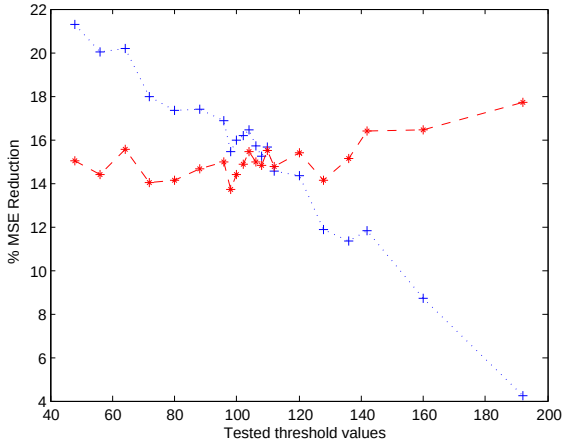

Figure 8: % MSE improvement in masked region of edge-evolved filter against db4 wavelet (dashed) and globally evolved filter (dotted).

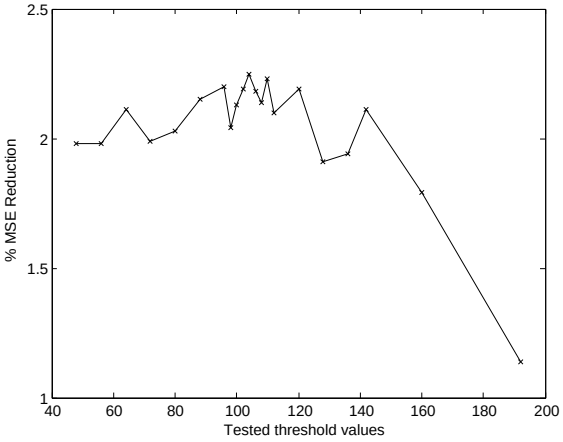

Figure 9: % MSE improvement for entire image using combined reconstruction over reconstruction with globally evolved filter only.

of 106.108 representing a reduction of 23.18%. Local filters for the reconstruction of object edges are evolved at various threshold settings ranging from 48 up to 192. Recall that during evolution, MSE fitness is only assessed near the object edges of the training image as indicated by the mask created at the given threshold. At each threshold, the locally evolved filter response is compared to the db4 DWT^{-1} response and the globally evolved filter response for the edge portions of the training image isolated by the mask.

The % reductions in MSE of the local filter against the global filter and the wavelet are plotted in figure 8. The local filters exhibit a reduction typically ranging from 14–16% against the db4 wavelet. Against the globally evolved filter, the local filters demonstrate an improvement of 21.30% at a conservative threshold of 48. The degree of improvement steadily declines as the threshold increases but remains significant at thresholds below 120. This makes sense because at low thresholds, only dark positions of the edge detection algorithm, indicating large intensity transitions between neighboring pixels (strong edges) are encompassed by the mask images. At higher thresholds, the masks are less selective and encompass larger portions of the image. The two plots cross at a threshold of 112; at this point the responses of the wavelet and the globally evolved filter are approximately equal. From this point, the wavelet outperforms the global filter in the mask-encompassed portion of the image to an increasing degree as the threshold decreases. The reverse is true as the threshold increases. This confirms that the globally evolved filter, while reducing error across the entire image, actually increases reconstruction error near object edges. However, the locally evolved filters provide consistent improvement near object edges.

Figure 9 plots the overall reduction in MSE with the image reconstructed using the combined locally and globally evolved filters versus the globally evolved filter alone at each tested threshold level. Improvement ranges from 1.98–2.25% at thresholds under 120, with the best performance coming at a threshold of 104. At this threshold, the MSE reduces from 106.108 to 103.719, a reduction of 2.25%. While this may not seem to be a significant improvement across the entire image, this improvement occurs strictly near object edges, such as building outlines or aircraft profiles. The portions of the image most critical for intelligence analysis demonstrate significantly improved reconstruction.

4.2 Performance on Satellite Image Set

The results above demonstrate trends in filter response, but because each is the result of a single GA run on a single image, statistically sound conclusions may not yet be drawn. In order to assess the performance of the proposed technique on a wider range of conditions, replicated GA experiments are conducted at multiple threshold levels for the four satellite images presented in figure 5. For each image, a global reconstruction filter is evolved, as well as an edge-isolated local filter using masks created at threshold levels 48, 88, and 104 in a single experimental replication. Fifteen total replications are conducted for each image, to enable statistical comparison of results. The responses of the wavelet, the globally evolved filter, and the locally evolved filter are recorded at each threshold level for the mask-enclosed portion of the image.

MSE of db4 Wavelet and Globally Evolved Filters

		global filter		% improvement	
image	wavelet	avg	std	avg	std
AF museum	151.25	120.22	1.11	20.51	0.71
Baghdad	194.08	159.89	0.23	17.62	0.12
B-52s	123.46	99.87	1.66	19.10	1.35
WPAFB	152.62	131.90	1.31	13.57	0.86

Table 1: MSE of images reconstructed with db4 wavelet and globally evolved filters.

Table 1 presents the MSE of each image reconstructed with the db4 $\mathrm{DWT}^{-}1$ and the average MSE achieved with the globally evolved filters across all replications. Evolved filters reduce MSE by between 13.57 and 20.51% on average, depending on the image. The masks created for each threshold encompass varying amounts of the training im-

% of image enclosed by mask

image	threshold 48	88	104
AF museum	5.738	9.349	11.275
Baghdad	24.664	36.481	41.557
B-52s	8.226	12.281	14.226
WPAFB	12.144	17.768	20.415

Table 2: % of images covered by masks created at each threshold.

% Improvement in mask region over wavelet using local filters

threshold	48		88		104	
image	avg	std	avg	std	avg	std
AF museum	10.90	0.36	**13.55**	**0.32**	**13.69**	**0.32**
Baghdad	**18.23**	**0.03**	18.2	0.03	18.00	0.11
B-52s	**22.51**	**0.16**	22.32	0.18	22.03	0.17
WPAFB	11.65	0.28	12.20	0.35	**12.55**	**0.32**

% Improvement in mask region over global filters using local filters

threshold	48		88		104	
image	avg	std	avg	std	avg	std
AF museum	**24.06**	**3.51**	20.36	3.11	18.7	3.14
Baghdad	**1.13**	**0.29**	0.78	0.25	0.59	0.28
B-52s	**13.97**	**4.31**	**14.04**	**5.05**	**12.85**	**3.92**
WPAFB	**11.10**	**2.63**	**9.48**	**2.18**	**8.75**	**2.24**

Table 3: Mean % MSE reduction in area enclosed by masks using combined filter reconstruction.

ages. Table 2 provides the % of each image enclosed by the masks for each threshold value. The masks for the Baghdad image encompass between 24.664 and 41.557% of the image; as seen in figure 7 (top right), the large number of structures in this image lead to a large number of edges in the image. The remaining images are much more edge sparse, with the Air Force museum image containing the smallest percent covered by the edge mask.

The locally evolved filters are compared within the mask-enclosed regions for each image in table 3. This table shows the average reduction of MSE over the wavelet (top) and over the evolved filters (bottom) for the given image. The best result and any results not statistically significantly different are shown in **bold** for each image. T-tests at significance level $\alpha = 0.05$ provides assessments of the differences between results. Within the mask-enclosed region, the local filters perform very well compared to the wavelet. The threshold value does not appear to strongly influence performance for the given images. The local filters demonstrate the best performance at a threshold of 104 for two images, and only slightly lower than the best performance for the two remaining images. Compared to the globally evolved filters, the local filters show significantly improved results for three

% Improvement in whole image over wavelet using combined filters

threshold	48		88		104	
image	avg	std	avg	std	avg	std
AF museum	1.00	0.03	1.59	0.04	**1.86**	**0.04**
Baghdad	5.75	0.01	8.18	0.01	**9.08**	**0.09**
B-52s	3.19	0.02	4.58	0.04	**5.13**	**0.04**
WPAFB	2.02	0.05	2.99	0.09	**3.43**	**0.18**

% Improvement in whole image over global filters using combined filters

threshold	48		88		104	
image	avg	std	avg	std	avg	std
AF museum	2.69	0.51	**3.28**	**0.62**	**3.41**	**0.69**
Baghdad	**0.36**	**0.09**	**0.35**	**0.12**	**0.31**	**0.12**
B-52s	2.23	0.77	**3.08**	**1.02**	**3.35**	**1.13**
WPAFB	**2.22**	**0.58**	**2.57**	**0.67**	**2.70**	**0.74**

Table 4: Mean % MSE reduction of entire images using combined filter reconstruction.

of four images, with the best performance typically coming at a threshold of 48, consistent with the plot in figure 8 demonstrating greater improvement at smaller threshold values. The locally evolved filters provide only minor improvement over the global filters for the Baghdad image. Recall from table 2 this image contains the greatest degree of edge transitions. Filters trained on this image exert a relatively large amount of selective pressure on filters providing improved reconstruction near object edges. Images containing fewer edges provide less evolutionary pressure, preferring improved reconstruction across the entire image at the expense of reconstruction near edges.

Table 4 shows MSE improvement when combining the global and local filters for reconstruction. The top portion shows improvement over the db4 DWT^{-1} when the mask-covered portion of the image is reconstructed with the local filter and the remainder with the wavelet. Because masks enclose a small portion of each image, improvement is modest. These small gains are measured across the entire image, though *all improvement comes near edges*. Not surprisingly, the greatest error reduction comes for the Baghdad image, containing the largest mask coverage. The best improvement comes at a mask threshold of 104 for each image, consistent with the trends illustrated in figure 9. The lower portion of table 4 provides results seen for reconstruction using both local and global evolved filters over use of globally-evolved filters alone. The Baghdad image demonstrates the smallest improvements; this image provides sufficient selection pressure for edge reconstruction as a training image. The remaining images show reconstruction improvement of between 2.22 and 3.41%, even though a relatively small portion of each image is covered by the mask. The best improvements typically occur at the 104 threshold, where a larger amount of image coverage provides greater room for improvement.

5. CONCLUSIONS AND FUTURE WORK

Existing techniques of filter evolution potentially provide significant improvement over standard wavelet transforms, but they often increase the error present near the edges of objects. Image processing applications such as target recognition and intelligence gathering cannot afford this resolution loss in the most critical sections of the image. The use of an edge detection algorithm and an edge-enclosing mask allows the evolution of reconstruction filters that improve the reconstruction of object edges by as much as 24% under conditions subject to high quantization error. Combined with filters evolved for entire images, the locally evolved filters provide a reconstruction algorithm suitable for applications requiring maximum object resolution while maintaining maximal compression ratios.

Results indicate that there may exist a correlation between the degree of edges within an image and the potential improvement a locally evolved filter may provide. Future experiments should focus on images containing a wide range of edge sparseness or abundance. Experiments should focus on the determination of appropriate mask creation thresholds for images of various edge abundance. These experiments will lead to the development of a system that given an image, determines the appropriate threshold setting and selection of appropriate filters from a library of previously evolved filters.

An additional interesting experiment may be to evolve filters trained using only the portions of the image containing no significant edges. By ignoring edges, such filters may demonstrate an improved response over globally evolved filters to the edge-sparse portions of images. Combined with edge-trained filters, such filters may provide further improvements in image reconstruction over wavelets as well as and filters trained to reconstruct an entire image.

6. REFERENCES

[1] A. Bruckmann, T. Schell, and A. Uhl. Evolving subband structures for wavelet packet based image compression using genetic algorithms with non-additive cost functions. In *Proceedings of the International Conference on Wavelets and Multiscale Methods*, 1998.

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

[3] G. Davis and A. Nosratinia. Wavelet-based image coding: an overview. *Applied and Computational Control, Signals, and Circuits*, 1(1), 1998.

[4] A. Gersho and M. Gray. *Vector Quantization and Signal Compression*. Kulwer Academic Publishers, 1991.

[5] Google. Google earth plus. http://earth.google.com/, 2006.

[6] U. Grasemann and R. Miikkulainen. Evolving wavelets using a coevolutionary genetic algorithm and lifting. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO-04*, volume 3103 of *Lecture Notes in Computer Science*, pages 969–980. Springer-Verlag, 2004.

[7] U. Grasemann and R. Miikkulainen. Effective image compression using evolved wavelets. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'05)*, pages 1961–1968, 2005.

[8] Y. Hill, S. O'Keefe, and D. Thiel. An investigation of wavelet design using genetic algorithms. In *Microelectronic Engineering Research Conference*, 2001.

[9] T. Hopper, C. M. Brislawn, and J. N. Bradley. Wsq gray-scale fingerprint image compression specification. Technical Report IAFIS-IC-0110, Federal Bureau of Investigation, February 1993.

[10] E. Jones, P. Runkle, N. Dasgupta, L. Couchman, and L. Carin. Genetic algorithm wavelet design for signal classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(8):890–895, August 2001.

[11] D. Kaur and L. Ying. Creating a neuro-fuzzy model by combining filtered images with various filtering operators for the detection of edges in new images. Technical report, University of Toledo, 2006.

[12] M. Lankhorst and M. van der Lann. Wavelet-based signal approximations with genetic algorithms. In *Proceedings of the 4th Annual Conference on Evolutionary Programming*, pages 237–255, 1995.

[13] F. Moore. A genetic algorithm for evolving improved mra transforms. *WSEAS Transactions on Signal Processing*, 1(1):97–104, 2005.

[14] F. Moore. A genetic algorithm for optimized reconstruction of quantized signals. In *IEEE Congress on Evolutionary Computation (CEC) Proceedings vol. 1*, pages 105–111, 2005.

[15] F. Moore, P. Marshall, and E. Balster. Evolved transforms for image reconstruction. In *IEEE Congress on Evolutionary Computation (CEC) Proceedings vol. 3*, pages 2310–2316, 2005.

[16] M. R. Peterson, G. B. Lamont, and F. Moore. Improved evolutioanry search for image reconstruction transforms. In *Proceedings of the IEEE World Congress on Computational Intelligence*, pages 9785–9792, 2006.

[17] B. S. Rani and S. Renganathan. Wavelet based texture classification with evolutionary clustering networks. In *TENCON 2003: IEEE Conference on Convergent Technologies for Asia-Pacific Region*, volume 1, pages 239–243, 2003.

[18] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1964.

[19] I. E. Sobel. *Camera Models and Machine Perception*. PhD thesis, Electrical Engineering Department, Stanford University, Stanford, CA, 1970.

[20] W. Sweldens. The lifting scheme: a custom-design construction of biorthogonal wavelets. *Journal of Aplied and Computational Harmonic Analysis*, 3(2):186–200, 1996.

[21] B. E. Usevitch. A tutorial on modern lossy wavelet image compression: foundations of jpeg 2000. *IEEE Signal Processing Magazine*, pages 22–35, September 2001.

[22] J. Walker. *A Primer on Wavelets and Their Scientific Applications*. CRC Press, 1999.

[23] A. H. Wright. Genetic algorithms for real parameter optimization. In G. Rawlins, editor, *Foundations of Genetic Agorithms*, pages 205–220, San Mateo, 1991. Morgan-Kaufman.