A Differential Evolution Algorithm for Optimizing Signal Compression and Reconstruction Transforms

Frank Moore and Brendan Babb
Mathematical Sciences Department
University of Alaska Anchorage
3211 Providence Dr., Anchorage, AK 99508
1-907-786-1742
brendanbabb@gmail.com, affwm@uaa.alaska.edu

Abstract. State-of-the-art image compression and reconstruction techniques utilize wavelets. Beginning in 2004, however, a team of researchers at Wright-Patterson Air Force Base (WPAFB), the University of Alaska Anchorage (UAA), and the Air Force Institute of Technology (AFIT) has 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. Unfortunately, the computational cost of our GA-based approach has been enormous, necessitating hundreds of hours of CPU time, even on supercomputers provided by the Arctic Region Supercomputer Center (ARSC). The purpose of this investigation was to begin to determine whether an alternative approach based upon differential evolution (DE) [20] could be used to (a) optimize transforms capable of outperforming those evolved by the GA, (b) reduce the amount of computation necessary to evolve such transforms, and/or (c) further reduce the mean squared error (MSE) of transforms previously evolved via our GA.

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, Satellite Images, Image Compression, Image Reconstruction.

1. INTRODUCTION

DE [20] is a fast, simple-to-use, yet powerful global optimization algorithm. DE is increasingly being used to solve a wide range of difficult real-world optimization problems, including ground state structure prediction for silicon-hydrogen cluster design; compressor supply system optimization; multi-sensor fusion; medical image registration; digital filter design; and analysis of X-ray reflectivity data (summarized in [20], pp. 311-511). For the research described in this paper, DE was applied to the optimization of real-valued coefficients describing transforms that outperform wavelets for the compression and subsequent reconstruction of images under conditions subject to quantization error.

Since the late 1980s, engineers, scientists, and mathematicians have used wavelets to solve a wide variety of difficult problems, including fingerprint compression [4], signal denoising [7], and medical image processing [1]. Adoption of the Joint Photographic Experts Group's JPEG2000 standard [22] has established wavelets as the principal methodology for image compression and reconstruction. JPEG2000 utilizes wavelets to improve upon the compression capabilities of previous JPEG [9] and JPEG-LS [10] standards.

Wavelets [6] 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 9/7 wavelet, these sets consist of the following floating-point coefficients (rounded to five decimal places):

h1 = [0.03783, -0.02385, -0.11062, 0.37740, 0.85270, 0.37740, -0.11062, -0.02385, 0.03783]
g1 = [0.06454, -0.04069, -0.41809, 0.78849, -0.41809, -0.04069, 0.06454]
h2 = [-0.06454, -0.04069, 0.41809, 0.78849, 0.41809, -0.04069, -0.06454]
g2 = [0.03783, 0.02385, -0.11062, -0.37740, 0.85270, -0.37740, -0.11062, 0.02385, 0.03783]

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 ([22], 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**

GECCO'08, July 12–16, 2008, Atlanta, Georgia, USA.

and $\mathbf{h^1}$, $\mathbf{d^1}$, and $\mathbf{v^1}$ are its first horizontal, diagonal, and vertical fluctuation subimages, respectively [23]. Using the multi-resolution analysis (MRA) scheme [11], a one-level DWT may be repeated $k \leq \log_2(\min(M, N))$ times. Note that the size of the trend signal $\mathbf{a^i}$ at level i of decomposition will be $1/4^i$ times the size of the original image $\mathbf{f}$ (e.g., a three-level transform produces a trend subimage $\mathbf{a^3}$ that is $1/64^{th}$ the size of $\mathbf{f}$). 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 $\mathbf{f}$.

Quantization [12] 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 $\mathbf{y}$ onto a smaller range of values $Q(\mathbf{y})$. 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 $\mathbf{y}$ 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 in the dequantized signal, $\rho(\mathbf{y}, \hat{\mathbf{y}})$, may be computed as the MSE for each sample.

For many applications, quantization is the most significant source of distortion in digital images. Unfortunately, the performance of wavelets degrades in proportion to the amount of quantization error; for critical applications of wavelets to military, security, or medical imaging tasks, such error may be unacceptable.

2. PREVIOUS RESEARCH

In a series of recent projects, researchers at the University of Alaska Anchorage, in cooperation with teams at the Air Force Institute of Technology, Wright State University, and Wright-Patterson Air Force Base, have investigated the possibility of using GAs to evolve sets of coefficients describing novel transforms that outperform wavelets for image processing applications subject to quantization error. The objective of these investigations has been to minimize the aggregate MSE present in reconstructed images, while producing transforms that match or exceed the compression capabilities of wavelets.

1. First, we showed that a GA [8] could be used to evolve coefficients describing an inverse transform capable of reducing the mean squared error (MSE) in reconstructed one-dimensional signals previously compressed by a DWT and subjected to quantization error. Results were promising [15], with error reductions consistently exceeding 91% for sinusoidal signals.
2. Next [19], we demonstrated that this approach could be successfully applied to photographic images. Our GA evolved inverse transforms capable of reducing MSE by as much as 10.7% in comparison to the selected wavelet.
3. Next [2], we 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 to the Daubechies-4 (D4) transform under conditions subject to a quantization step of 64, while maintaining an average compressed file size (FS) less than or equal to the FS produced by the D4 wavelet.
4. Next [17], we utilized the massive computational power of supercomputers at the Arctic Regional Supercomputer Center (ARSC) to evolve one-level transforms. For a quantization step of 64, these transforms reduced MSE by nearly 40% (2.203 dB) for the training image, and by an average of nearly 23% (1.126 dB) on test images. In addition, according to an Information Entropy (IE) measure commonly used to accurately estimate FS, the average compressed FS for evolved transforms was less than or equal to that of the D4 wavelet.
5. Next [18], we used a GA to evolve three-level MRA transforms described by a single set of coefficients used at every level. The resulting transforms were capable of an average MSE reduction of 7.61% (0.34 dB) under conditions subject to a quantization step of 64, while keeping FS in check.
6. Next [16], we expanded our GA to evolve three-level MRA transforms that utilized a different set of coefficients at each MRA level. Each individual consisted of 48 real-valued coefficients (16 for each MRA level). At quantization equal to 64, the evolved MRA transform reduced MSE by as much as 12.92% (0.60 dB), again while keeping average FS less than or equal to the FS produced by the three-level D4 MRA transform.
7. Finally [3], we demonstrated that our approach could be used to evolve four-level MRA transforms that outperformed the biorthogonal 9/7 wavelet filter pair incorporated in the Federal Bureau of Investigation (FBI) fingerprint compression standard [5]. Each individual used the 9/7 structure at each MRA level, and thus consisted of 128 real-valued coefficients (16 forward and 16 inverse transform coefficients at each level). For each compressed image, only the 6.25% largest values were retained, i.e., images were subjected to a 16:1 quantization. These tests produced the following results:
   a) The best transform evolved by the GA reduced MSE by an average of 24.03% (1.20 dB) on the four fingerprint images used for training.
   b) The best transform averaged 16.01% (0.76 dB) MSE reduction when subsequently tested against a population of 80 fingerprint images.
   c) The average size FS compressed by the evolved transform was virtually identical to the FS produced by the 9/7 wavelet.
   d) Evolved transforms were subsequently tested on photographs commonly used by the signal processing community, such as “zelda”, “lenna”, and “airplane”. The MSE of the evolved transforms was consistently worse on these images than the original 9/7 wavelet. This result suggests that the GA is capable of automatically discovering and exploiting specific features of fingerprints that do not commonly appear in other photographic images.

For the first five tasks, the GA seeded each individual in 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. For the final two tasks, the coefficients at each level of the transform were independently initialized to a different randomly mutated copy of the selected wavelet’s coefficients.

The published research most closely related to this project combined a coevolutionary GA [13] with the lifting scheme [21] to evolve wavelets specifically for fingerprint images. The best solutions evolved by those researchers "averages 0.75 dB quality improvement over the FBI wavelet" when subsequently tested on a population of 80 fingerprints [14]. Thus, our best evolved transform approximately equaled the performance of their best wavelet. However, it should be noted that our GA was not constrained to produce transforms having the precise mathematical properties of wavelets [4], such as biorthogonality. Instead, our GA was free to evolve whatever combination of wavelet and scaling coefficients resulted in the most effective MSE reduction. This additional freedom allowed our approach to more effectively search the space of both wavelets and non-wavelet transforms in order to better compensate for quantization error.

3. THE DIFFERENTIAL EVOLUTION ALGORITHM

The goal of a typical 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 compressed file size *FS*).
2. The distortion observed in the reconstructed image produced by the corresponding inverse transform (i.e., the *SE*).

For the tests described in this paper, unlike our previous research, *FS* was determined by the selected quantization scheme. Using the technique employed in similar research projects [14] to maximize comparability of results, all but the largest 6.25% of the compressed image's values were discarded (i.e., our approach retained only 1/16th of these values). In effect, then, the fitness of solutions evolved by our DE-based approach could be evaluated solely by comparing the original images from the training set to the images produced after compression via the evolved forward transform, quantization via the technique stated above, encoding, decoding, dequantization, and reconstruction via the corresponding evolved inverse transform.

The training images used during this research were four fingerprint images selected from a standard training set of 80 fingerprints. Our previous research [3] using a GA had demonstrated that the use of fingerprints with higher energy content (i.e., "sharper" images) contributed to the evolution of better transforms. The same fingerprints used by the GA for previous research were used by our DE algorithm to allow direction comparison with our previous results.

4. TASK 1: CAN DE ALONE EVOLVE POWERFUL NEW TRANSFORMS?

The first task accomplished during this research determined the utility of evolving novel compression and reconstruction transforms using DE alone. A major step towards completion of this task involved setting up five test computers with Matlab and DE software at the University of Alaska Anchorage (UAA). Only after preliminary tests were satisfactorily completed was it possible to launch the large-scale training runs necessary to collect data for this study.

DE is notoriously sensitive to control parameters NP (the size of the evolving population) and G (the number of times DE traverses the population); for this reason, five tests were run using different combinations of these parameters, as follows:

Test	NP	G
1	100	5000
2	50	10000
3	25	20000
4	200	2500
5	150	3333

Note that each test evaluated approximately 500,000 individuals (for example, test 2 evaluated a population of 50 individuals over 10,000 generations).

For each of these runs, the initial population was seeded with one exact copy and NP-1 mutated copies of coefficients from the 9/7 wavelet at each of four MRA levels. The training population consisted of four fingerprint images from a data set provided by the University of Texas at Austin. The selected quantization technique retained the largest 6.25% of the values from the compressed signal (a technique employed in our previous GA research). The fitness of each individual was computed as the mean squared error (MSE) of the reconstructed fingerprint images in comparison to the training set. Each test used DE strategy 5 (DE/rand/1 with per-generation dither).

The results of this task are shown in Fig. 1. These results substantiate the following claims:

1. DE can produce transforms that outperform the 9/7 wavelet for fingerprint compression and reconstruction under conditions subject to quantization. The best DE training run completed during the research period described in this report produced coefficients defining a transform that reduced the MSE in reconstructed fingerprints by nearly 0.5 dB.
2. When using the 9/7 transform to create the initial population, the success of DE is highly dependent upon the choice of control parameters NP and G. Test 2 evolved a better transform than test 3, which evolved half as many candidate solutions over twice the number of generations. Test 2 also evolved a better transform than tests 4 and 5, which used larger populations and fewer generations. (NOTE: test 1 crashed prior to completion, as indicated by the *dnf* entry in the table shown above; thus, it is currently unknown whether NP = 100 is a better choice than NP = 50 for this problem. Since each run required approximately 20 days to complete, rerunning test 1 to completion was infeasible.)
   3. Each of the approximately 500,000 candidate solutions was applied to a training population of four fingerprints; thus, each run required the compression, quantization, encoding decoding, dequantization, and reconstruction of approximately 2,000,000 fingerprints. While this number may seem enormous, it must be pointed out that the GA developed during our previous research required several times that amount of computation to produce the state-of-the-art transforms described in our previous research, which achieved 1.2 dB reduction on the same training set. Whether or not our DE-based approach could achieve a 1.2 dB MSE reduction more rapidly than the GA remains an open question.

Transform	MSE	Reduction
9/7	63.613500	--
DE-optimized from 9/7 coefficients (test 1)	*- dnf -*	*??*
DE-optimized from 9/7 coefficients (test 2)	57.052467	10.31%
DE-optimized from 9/7 coefficients (test 3)	63.609533	0.01%
DE-optimized from 9/7 coefficients (test 4)	61.321511	3.60%
DE-optimized from 9/7 coefficients (test 5)	60.697778	4.58%

Fig. 1. Results from Various Training Runs of the DE System

Likewise, whether or not DE could achieve more than 1.2 dB reduction in the same amount of computation also remains an open question. The answers to these questions will necessitate many months of additional processing time, and thus are beyond the scope of the current project. However, the fact that only 2,000,000 evaluations had already produced transforms capable of 0.5 dB MSE reduction is quite promising.

5. TASK 2: CAN DE FURTHER IMPROVE THE PERFORMANCE OF TRANSFORMS PREVIOUSLY EVOLVED VIA A GENETIC ALGORITHM?

The second task accomplished during this research determined whether the compression and reconstruction capabilities of transforms previously optimized by our genetic algorithm (GA) could be further enhanced via DE.

Five training runs were performed using NP and G values that were identical to those used for Task 1; thus, each run evaluated approximately 500,000 individuals.

For each run, the initial population was seeded with one exact copy and NP-1 mutated copies of a best-of-run solution previously produced by an extended GA run on *nelchina*, a powerful parallel processor only recently decommissioned by the Arctic Region Supercomputer Center (ARSC). As with Task 1, the training population consisted of four fingerprint images from the University of Texas at Austin set. Compressed fingerprints were subjected to 16:1 quantization, and MSE was again used as a measure of fitness. As with Task 1, each test used DE strategy 5 (DE/rand/1 with per-generation dither).

The results of this task are shown in Fig. 2.

Transform	MSE	MSE reduction
9/7	63.613500	--
GA-evolved	53.355822	16.125%
DE-optimized (test 1)	53.353756	16.128%
DE-optimized (test 2)	53.353567	16.129%
DE-optimized (test 3)	53.352978	16.129%
DE-optimized (test 4)	53.353278	16.129%
DE-optimized (test 5)	53.353755	16.128%

Fig. 2. Performance of Best-of-Run Solutions Produced by Starting from a Previously GA-Optimized Transform

These results appear to substantiate the following claims:

1. The amount of additional performance enhancement produced by applying DE to a transform previously evolved via our GA was negligible. Even after 500,000 iterations, our best DE run produced coefficients that improved performance by less than 0.005%. There are two possible explanations for this result: either
   (a) The GA had already produced a transform capable of near-optimal performance, thus leaving DE very little room for additional enhancement; or
   (b) The DE approach simply lacked the power necessary to effectively search the solution space for this problem, which requires the simultaneous optimization of 128 real-valued coefficients (sixteen coefficients for the compression transform and sixteen coefficients for the reconstruction transform at each of four MRA levels).
2. For Task 2, and for the selected optimization strategy (DE strategy 5), test results appear to indicate that the choice of NP and G were ultimately of little importance. Differences in the performance improvement produced by each of the five DE tests were miniscule. This outcome was surprising and contradicts our previous experience with DE, which typically indicate that the outcome of each run is highly sensitive to the choice of control parameters.

6. CONCLUSIONS

This research demonstrated each of the following key points:

1. DE is capable of evolving a set of 128 real-valued coefficients describing a four-level MRA transform that outperforms the 9/7 wavelet for fingerprint compression and reconstruction under conditions subject to quantization. In light of the enormous solution space of this problem, this result is quite impressive. It clearly demonstrates the potential of DE as a means for pushing the state-of-the-art in this important research area.
2. DE was incapable of subsequently optimizing a state-of-the-art transform previously evolved by our GA.

Statistical validation of the preliminary results described above will necessitate completion of a much greater number of training runs. Such tests will necessitate several weeks of computation. The relatively short duration of the Air Force Research Laboratory's Summer 2007 Extension Grant made completion of such tests infeasible.

7. FUTURE DIRECTIONS

Future research should address each of the following tasks:

(a) Given the same amount of computation, can DE evolve a better transform than our GA? Similarly, can DE evolve an equally capable transform using substantially fewer CPU resources? To answer these questions, we will need to utilize supercomputers such as those provided by the Arctic

Regional Supercomputer Center to complete several runs of similar magnitude to our previous GA runs. Only then can we qualitatively compare the transforms produced via both approaches.

(b) All of the tests described above used DE strategy 5. It is possible that other strategies may be more suitable for this problem. Future investigations should compare each strategy to determine whether any other approach works more effectively for this problem.

(c) Fingerprint compression is one of several important applications that currently utilize wavelets. Other areas include satellite images, medical images, and digital photography. Future research should optimize transforms for each of these applications. Careful study of these transforms may reveal why they are more effective at compressing and reconstructing a particular signal class.

8. ACKNOWLEDGEMENTS

The research described in this paper was performed during the Summer 2007 term under an Extension Grant from the Air Force Research Laboratory. The author especially recognizes Pat Marshall (AFRL/SNAT) at WPAFB, Dayton, OH, for his continuing support. Brendan Babb has provided many hours of outstanding effort to push this technology forward. Our colleagues at Wright State University, the Air Force Institute of Technology, and WPAFB -- Michael Peterson, Gary Lamont, and Eric Balster -- have provided considerable assistance as this series of projects have evolved over the past four years. Special thanks are also due to John Graniero (JAG Technologies), Cynthia Cooley (AFRL/IFB), and Fran Connors (SUNY-IT) for their assistance with administrative details.

9. REFERENCES

[1] Aldroubi, A. and M. Unser (eds.) 1996. Wavelets in Medicine and Biology, CRC Press.

[2] 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.

[3] Babb, B. and F. Moore 2007. The Best Fingerprint Compression Standard Yet, Proceedings of the 2007 IEEE International Conference on Systems, Man, and Cybernetics, 10/7-10, 2007, Montreal, Quebec, Canada, IEEE.

[4] 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.

[5] Cohen, A., I. Daubechies, and J.-C. Feauveau 1992. Biorthogonal Bases of Compactly Supported Wavelets, Communications on Pure and Applied Mathematics, 45 (5): 485-560, June 1992.

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

[7] Donoho, D. 1993. Nonlinear Wavelet Methods for Recovery of Signals, Densities, and Spectra from Indirect and Noisy Data, Different Perspectives on Wavelets, American Mathematical Society.

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

[9] ISO 1994. ISO/IEC 10918-1 and ITU-T Recommendation T.81, Information Technology – Digital Compression and Coding of Continuous-tone Still Images: Requirements and Guidelines.

[10] ISO 1999. ISO/IEC 14495-1 and ITU-T Recommendation T.88, Information Technology – Lossless and Near-lossless Compression of Continuous-tone Still Images.

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

[12] Mallat, S. 1998. A Wavelet Tour of Signal Processing, Academic Press.

[13] Grasemann, U. and R. Mikkulainen 2004. Evolving Wavelets Using a Coevolutionary Genetic Algorithm and Lifting, Proceedings of the Sixth Genetic and Evolutionary Computation Conference (GECCO'2004), II: 969-980, Springer.

[14] Grasemann, U. and R. Mikkulainen 2005. Effective Image Compression using Evolved Wavelets, Proceedings of the Seventh Annual Genetic and Evolutionary Computation Conference (GECCO'2005), 6/25-29, 2005, Washington, DC, 2: 1961-1968, ACM.

[15] Moore, F. 2005. A Genetic Algorithm for Optimized Reconstruction of Quantized Signals, Proceedings, 2005 IEEE Congress on Evolutionary Computation, IEEE.

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

[17] Moore, F. and B. Babb 2006. Revolutionary Image Compression and Reconstruction via Evolutionary Computation, WSEAS Transactions on Signal Processing, 2(9): 1203-1208, WSEAS.

[18] Moore F. and B. Babb 2006. Revolutionary Image Compression and Reconstruction via Evolutionary Computation, Part 2: Multiresolution Analysis Transforms, WSEAS Transactions on Signal Processing, 2(9): 1209-1214, WSEAS.

[19] Moore, F., P. Marshall, and E. Balster 2005. Evolved Transforms for Image Reconstruction, Proceedings, 2005 IEEE Congress on Evolutionary Computation, IEEE.

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

[21] Sweldens, W. 1996. The Lifting Scheme: A Construction of Biorthogonal Wavelets, Journal of Applied and Computational Harmonic Analysis, 3(2): 186-200.

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

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