

New Perspectives on Image Compression

Michael Thierschmann, Reinhard Köhn, Uwe-Erik Martin
LuRaTech GmbH
Berlin, Germany

Abstract

Effective Data compression techniques are necessary to deal with the increasing amount of data produced by modern image processes. The quality of the currently established image compression techniques, including the JPEG-standard, are not sufficient for an increasingly large number of applications. And with the increasing number of network based online applications, additional features like quick look generation, progressive transmission, and scalability in quality and resolution of the image are new requirements.

New image compression techniques based on wavelet-transformations offer a solution. These techniques allow the combination of lossless and lossy compression with extremely high compression rates and image quality. This paper discusses the principles of wavelet- and JPEG-compression and compares the new wavelet techniques with the JPEG-standard method. Additionally we present a software image compression package, LuraWave, which allows the use of wavelet image compression in a wide range of commercial applications.

Introduction

The increasing amount of data produced by modern image generating and processing systems is often limited by storage and transmission capacity. Effective data compression techniques are necessary to overcome this problem. The JPEG-standard is one such compression technique.

The JPEG-standard is a powerful image compression technique providing lossy image compression with a high visual quality. Nevertheless, there are applications which require a higher image quality than JPEG can provide. In very demanding applications the JPEG technique has significant problems, for example increasing compression rates result in rapidly degrading image quality and annoying blocking artefacts. The completely different techniques used by JPEG to handle lossless and lossy compression may also pose problems for some applications.

New techniques based on wavelet transformations offer a solution, allowing lossless compression and very high compression rates with high image quality. The artefacts introduced at high compression rates are of a very different structure to those introduced by the JPEG technique and can be influenced by a wide range of wavelet filters.

Even in areas where JPEG performs optimally, wavelet compression delivers better image quality without requiring significantly more computational resources. Furthermore, the increased functionality of this modern compression technique opens up completely new possibilities in application design and operation.

This paper discusses the principles of wavelet- and JPEG compression and compares the new wavelet techniques with the JPEG standard method. This comparison is illustrated by visual examples. Additionally we present the commercial wavelet compression software LuraWave as an example of a practical implementation of wavelet image data compression.

Image classes

In choosing a suitable image compression technique one has to consider not only the type of compression required, for example lossless or lossy, but also characteristics of the images to be compressed.

We can divide images into two classes: images from a natural source, for example photographs and scans with a high color resolution, we shall refer to as "natural images"; artificially generated images, for example line drawings and cartoon-like images, we shall refer to as "artificial images".

Images of these classes have distinct characteristics. Natural images tend to show a continuous distribution of colors and seldom abrupt changes of color at object edges. In contrast artificial images often have a discontinuous color distribution and abrupt edges without transitions. Currently there is no compression technique which can handle both types of images, they both need special adapted compression algorithms. For this special purpose the company LuRaTech has developed an algorithm, which automatically determines the image class and uses the optimal compression method. It can be operated on mixed images, too. In this case it separates the original into subimages, which are compressed independently and recombined after decompression. This technology is available in the LuraDocument product line.

Overview of image compression techniques

Currently used image compression techniques can be separated into two groups: lossless and lossy compression. Lossless compression allows the reconstruction of the original image data from the compressed image data. With lossy compression a higher compression rate is possible by allowing small differences between original and reconstructed images.

Comression rates for natural images with lossless compression are generally small, ranging from 1:1 for uncompressable data to 1:3. Typical values, using TIFF-LZW encoding for example, are around 1:1,5.

Lossy compression achieves compression rates of 1:5 to 1:30 using standard techniques (e.g. JPEG) and up to 1:300 using newer techniques (e.g. wavelet, fractal compression). Most lossy compression techniques are based on two-dimensional transforms, followed by quantisation and encoding stages. The loss of information is introduced by the quantization stage which intentionally rejects less relevant parts of the image information (disregarding rounding errors in the transformation step).

Specially adapted techniques which achieve high compression rates exist for artificial images. These techniques are based on run length coding, sometimes followed by an entropy coder. Examples of this techniques are the Fax-encoding standard and the image formats PCX and RLE.

The confusingly large number of compression techniques prevents many potential users from directly applying image compression. A technique and an image format which covers a wide range of alternative compression requirements would be preferable. JPEG is one attempt at creating such a technique for natural images, assembling a number of compression techniques in a common standard. The following section describes its functionality.

While JPEG integrates a wide range of compression techniques, these methods can only be realized by introducing several very different working modes. In contrast wavelet compression allows the integration of various compression techniques into one algorithm.

The JPEG Standard

The popular current image compression techniques are implementations of the JPEG-standard (ITU T.81). The JPEG-standard covers a collection of compression techniques: it contains a lossless mode; two progressive modes; and a simple lossy technique, which encodes the image sequentially from the upper left to the lower right corner. Every mode can be split into different versions, so for instance one can chose between Huffman and arithmetic coding, though some useful combinations are partially missed, for example JPEG does not contain a progressive lossless mode of operation.

The large number of different coding techniques makes a complete implementation of the JPEG-standard extensive, while a minimal JPEG decoder has only to realize the sequential lossy mode of operation. So if one has to choose the best suited JPEG compression technique one has the problem of ensuring that the target system can decode the images.

The lossy JPEG compression initially divides the source image into squares of 8x8 pixels. As the squares are then separately encoded, the reconstructed image may contain discontinuities at the boundaries of the squares, resulting in one of the major JPEG disadvantages. In addition the analyzing functions of the DCT are applied independent of its granularity on small image areas of the same sizes. Better results could be achieved by applying smoothing low pass functions to larger areas and differentiating high pass functions to smaller areas.

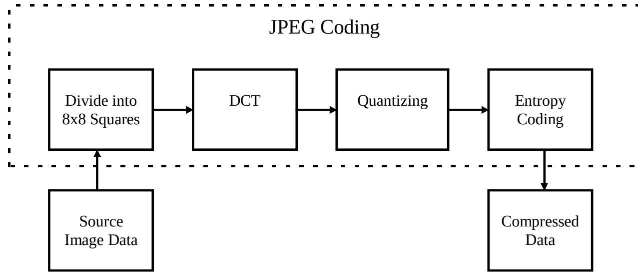

Figure 1 Simplified JPEG Coding Scheme for lossy compression

These blocking artefacts of the JPEG technique are visible only in enlarged images, but cause problems in many image processing applications. Particular problems occur when analyzing algorithms are applied to JPEG compressed images or when the images are re-rastered for printing.

When a compressed image is re-compressed (e.g., with a higher compression rate or after some processing) these additional edges further degrade the quality of the image.

Description of Wavelet compression

Wavelet compression applies a two-dimensional transformation to decorrelate the image information. Image compression with wavelet algorithms applies the steps image preparation, transformation, quantization and encoding. The quantization and encoding steps are usually combined allowing embedded coding, which will be explained later.

Image preparation

Wavelet transformation is applied recursively, dividing the original image matrix into 4 transformed matrices with half number of rows and columns. To ensure this division by two in every stage, the original image is supplemented by rows and columns to ensure the width and height are a multiple of 2^{level}.

A division into image blocks, as used by JPEG and fractal compression, is not necessary. Nevertheless, it can be applied to gain more efficient memory usage if required. If division is applied, the transformation is performed overlapping the block edges, so these edges will not influence the transformed image. The image is always transformed as a logical unit.

Transformation

The discrete wavelet transformation (DWT) is very similar to the dicrete cosine transformation (DCT), but it does not use the spatially unlimited sine and cosine functions for analysing the image. The base functions of the DWT are scaling and wavelet functions. These functions combine orthogonality, which is the first condition for transforming and identical reconstructing signals, with "compact support", i.e., they have a finite spatiality. This enables the analysis of signals without the windowing effects which result from the application of functions of infinite length to signals of finite length. For this reason images need not be split into blocks.

With the help of the most simple wavelet, the Haar-wavelet, we can explain the discrete wavelet transformation:

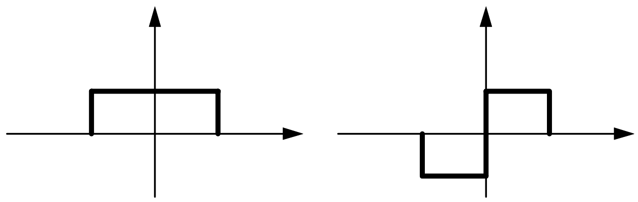

Figure 2: Basic functions for Haar transformation

In this simple case analyzing the signal with the Haar wavelet and scaling functions calculates the mean and the difference of two neighbouring pixels. The results are downsampled and stored as lowpass and highpass images. The lowpass image is now further analyzed in the same way.

The final transformed image consists of a number of highpass images of different sizes and a single lowpass image.

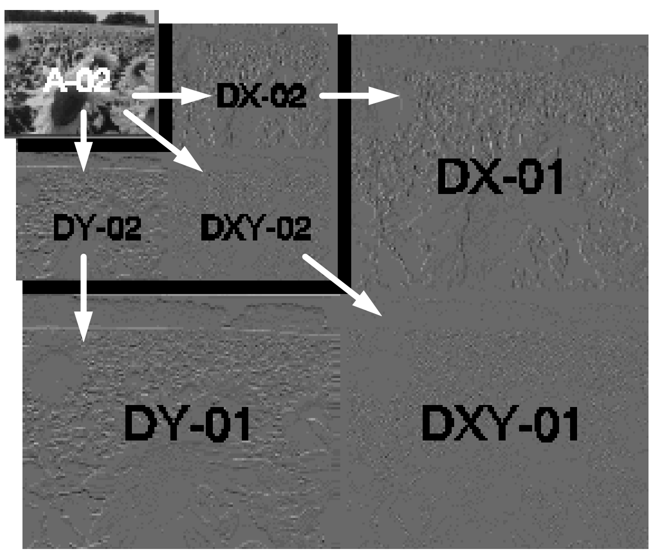

Figure 3: Example of a 2-Level, 2-dimensional Wavelet Transformation

The transform is normally applied with more function pairs than the Haar wavelet and scaling, e.g. Daubechies wavelet or biorthogonal spline wavelets.

The DWT can be calculated using either integer or floating point arithmetic. Floating point calculations allow a larger number of possible wavelet filters and the potential of high quality lossy compression. On the other hand, it has the disadvantage of rounding errors which need to be eliminated for lossless compression, resulting in higher expenses and calculation time. If one wants to compress images losslessly, the transform must be completely reversable, no rounding errors are allowed. This can be achieved by applying a special integer transformation.

Quantizing and coding

Quantization and encoding have to select the data being transmitted and put them into a suitable order in the data stream. They can be implemented as separate stages, as in JPEG, with the quantization stage rejecting some of the less relevant information and passing on the rest to the encoder stage for encoding.

By combining quantization and encoding into a single stage one gains the ability to exactly control the compression quality (image quality vs. Compression rate). The size of the encoded data stream can be predefined exactly. In this case quantization of the transformed coefficients occures during encoding.

The initial quantization is coarse with only the largest coefficients considered and encoded. The quantization intervals are then split and new, smaller coefficients, which were not considered in the first cycle, become significant. The values of the coefficients selected in the first cycle are specified more exactly by new information. The cycles continue until the data stream has reached its predefined length or until all information is encoded and the image is compressed losslessly. Applied to binary numbers and intervals, which are halfed in every cycle, the individual bit planes of the binary representation of the coefficients are encoded separately. This process is called bitplane coding. The resulting data stream is embedded, that is, the image information is ordered with respect to its importance in reconstructing the image. This allows the the reconstruction of a complete image from the initial part of a data stream, though of course with a reduction in quality when compared to the image which can be reconstructed from the whole stream. This feature has many potential advantages for applications in a wide range of fields, e.g. Internet and World Wide Web.

So wavelets can be used to form the basis of a compression technique which combines lossy and lossless compression. From lossless compressed image data one can obtain a higher compression rate simply by cutting the data stream. In addition the reconstructed images are of a higher quality than JPEG compressed images.

A comparision of JPEG and wavelet compressed images is shown in Figure 4 to Figure 6. The wavelet algorithm used is LuraWave.

Functional Advantages of Wavelet Compression

This section presents the functional advantages of wavelet compression and some of the resulting possibilities for application development.

Wavelet compression combines lossless and very high quality lossy compression in one algorithm. If an embedded coding scheme is used, lossless and lossy compression are even combined in one file! By using only a fraction of the lossless compressed data, it is possible to generate images with any higher compression rate. Although some quality will be lost, a significant reduction in transmission and decompression time is achieved. Applications which have to supply images with different resolutions and qualities (i.e. thumbnails, quick looks, full image, detail) need a very high administrational overhead if old compression techniques are

Figure 4: "Goldhill" original, 8 bit/pixel

Figure 5: "Goldhill" 1: 50, 0.16 bit/pixel, JPEG compressed

Figure 6: "Goldhill" 1:50, 0.16 bit/pixel, wavelet compressed with LuraWave

used, because each image needs to be generated, stored, and updated in several resolutions. Using wavelet compression, from one embedded compressed image, generated with the highest required quality, every smaller resolution can be generated without increased transmission times and with no computational overhead. This substantially decreases the administrational and storage costs for such applications.

Another advantage of wavelet compression is the absence of blocking artefacts, which are typical in JPEG compressed images. Highly wavelet compressed images have artefacts as well, but they are of a different structure with a direct relation to the image content. Discontinuities at block edges are avoided completely. This makes wavelet compressed images less problematic for analyzing applications. Studies showed, that three dimensional elevation models work well on images compressed 1:15 and higher [1].

Wavelet compression has short and nearly symmetric compression and decompression times. The complexity and the calculation times for wavelet compression are comparable to those for JPEG compression, despite the greater functionality of wavelet compression. They are well below that required for fractal compression. This makes wavelet compression suitable for realtime processing applications (i.e. scanners, digital cameras) and large images as well as for fast multimedia applications.

These advantages were realized by the JPEG group, too. The next JPEG standard, available in the beginning of the next century, will incorporate at least some of these new features. For developers, who want to design their products today with the advantages of future technology, the LuraWave product line is available. It features all descibed advantages using one of the most up to date wavelet compression algorithms. It is available in form of developer kits for nearly all operating systems, and DSP based hardware solutions are under development. All products will support the next generation JPEG standard as soon, as it is available. This way you can benefit from the wavelet technology today, without loosing the advantage of a standard in the future.

A range of applications which are based on the LuraWave wavelet compression is also available.

Compressing documents with bitonal areas and natural image areas

Wavelet compression offers superior performance on all natural images. On the other hand, for bitonal images wavelet compression produces unwanted artefacts at the extremely sharp edges. This is especially annoying in scanned documents, where full color images and text are combined.

To solve this problem, a separation algorithm was developed. This algorithm splits the original image in bitonal images, which will be compressed with lossless bitonal wavelet compression, preserving all edges perfectly, and in natural images, which are compressed using the wavelet compression discussed in this paper.

This solution still works, if text areas and image areas are overlapping. That way even complex documents like maps can be compressed with extremely high quality.

With the LuraDocument products, the company LuRaTech has integrated ist specialysed seperation algorithm with a modern wavelet image compression for natural images and a unique bitonal wavelet compression. Even pure bitonal images can be compressed with higher compression rates than the standard CCITT Fax Group 4 coder achieves. With just a little more compressed data, the compressed documents contain a full color representation of the original.

The LuraDocument software is available as a Photoshop PlugIn, a C-SDK and a Command Line Tool.

Providing wavelet compression for your application: LuraWave

Even the best compression algorithm is of no use, if no easy implemntation into applications is available. The LuraWave software solves this problem by delivering plug-ins for popular standard software (Adobe Photoshop, Macromedia Director, Netscape Navigator and MS Internet Explorer), stand alone viewers (LuraWave Studio), easy to use software development kits (Active X control for MS Windwos 95/NT), and a C library, which is available for nearly all operating systems (including most UNIX systems). Additionally there are Command Line Tools for batch processing on the most popular operating systems available. All products work on an interchangable file format. For computationally demanding applications, hardware based solutions will be available in the near future.

Summary

JPEG is a powerful image compression technique comprised of a collection of compression algorithms and modes of operation. Nevertheless there are many applications which require higher image quality than JPEG can provide. In practice many implementations of the JPEG standard use only the simplest mode of operation, with many of the other modes rarely used. The advantage of standardization is not realized for these modes.

Currently, in therms of compression quality, there exist better techniques than JPEG for image compression. Wavelet compression is a new and very efficient image compression technique, providing higher functionality and image quality. Lossless and lossy compression of natural images can be achieved with the same algorithm. The software package LuraWave supplies the functionality of wavelet compression to a wide range of commercial applications.

For documents with bitonal areas advanced separation algorithms are available. The company LuRaTech offers this innovative technology as ready to use software components in the LuraDocument product line. For more informations about wavelet compression have a look at our homepage http://www.luratech.com.

References

1. M. Thierschmann, P. Reinartz, M.Lehner, Suitability of Wavelet-Image-Data-Compression for the Derivation of Digital Elevation Models, SPIE Conference Wavelet Applications III, Vol. 2762, pp. 587-600, Orlando, Florida, 8th-12th April 1996
2. S. Mallat, A Theory for Multiresolution Signal Decompression, The Wavelet Representation, IEEE Trans. On Pattern Analysis and Machine Intelligence, No. 11, pp. 674-693, July 1989
3. J.M. Shapiro, Embedded Image Coding Using Zerotrees of Wavelet Coefficients, IEEE Trans. Signal Processing, Vol. 41, No. 12 pp. 3445-3462, 1993
4. A.K. Louis, P. Maass, A. Rieder, Wavelets: Theorie und Anwendungen, Teubner, Stuttgart, Germany, 1994