A Comparison of Lossless Image Compression Algorithms for Colour Retina Images

Gerald Schaefer
Aston University, UK

Roman Starosolski
Silesian University of Technology, Poland

INTRODUCTION

Diabetic retinopathy is the leading cause of blindness in the adult population. In order to effectively identify patients suffering from the disease, mass-screening efforts are underway during which digital images of the retina are captured and then assessed by an ophthalmologist. In order to identify features such as exudates and microaneurysms, which are typically very small in extent, retinal images are captured at high resolutions. This in turn means large file sizes and, considering the archival of typically thousands of records, a high demand on computational resources, in particular storage space as well as bandwidth when used in a Picture Archiving and Communications System (PACS). Image compression therefore seems a necessary step. Image compression algorithms can be divided into two groups: lossy techniques where some of the visually less important image data is discarded in order to improve compression ratios, and lossless methods which allow the restoration of the original data. As the features that indicate retinopathy are very small in size and following legislation in several countries, only lossless compression seems suitable for retinal images.

In this article, we present experiments aimed to identify a suitable compression algorithm for colour retina images (Schaefer & Starosolski, 2006). Such an algorithm, in order to prove useful in a real-life PACS, should not only reduce the file size of the images significantly but also has to be fast enough, both for compression and decompression. Furthermore, it should be covered by international standards such as ISO standards and, in particular for medical imaging, the Digital Imaging and Communication in Medicine (DICOM) standard (Mildenberger, Eichelberg, & Martin, 2002; National Electrical Manufacturers Association, 2004). For our study, we therefore selected those compression algorithms that are supported in DICOM, namely TIFF PackBits (Adobe Systems Inc., 1995), Lossless JPEG (Langdon, Gulati, & Seiler, 1992), JPEG-LS (ISO/IEC, 1999), and JPEG2000 (ISO/IEC, 2002). For comparison, we also included CALIC (Wu, 1997), which is often employed for benchmarking compression algorithms. All algorithms were evaluated in terms of compression ratio which describes the reduction of file size and speed. For speed, we consider both the time it takes to encode an image (compression speed) and to decode (decompression speed), as both are relevant within a PACS.

Experiments were performed on a large dataset of more than 800 colour retinal images, which were also divided into subgroups according to retinal region (nasal, posterior, and temporal) and images size. Overall, JPEG-LS was found to be the best performing algorithm as it provides good compression ratios coupled with high speed.

LOSSLESS IMAGE COMPRESSION ALGORITHMS

In this section, we give a brief overview of the lossless compression algorithms that we have evaluated. For further details on the algorithms, we refer the reader to the original references.

- **TIFF PackBits:** A simple RLE compression algorithm (Adobe Systems Inc., 1995). The Tag Image File Format (TIFF) standard specifies this simple runlength (RLE) coding technique; we used the LibTIFF implementation by Leffler (version 3.6.1, http://www.remotesensing.org/libtiff/).
- **Lossless JPEG:** Former JPEG (Joint Photographic Experts Group) committee standard for lossless image compression (Langdon et al., 1992). The standard describes predictive image

compression algorithm with Huffman or arithmetic entropy coder. We used the Cornell University implementation (version 1.0, ftp://ftp.cs.cornell.edu/pub/multimed/ljpg.tar.Z) which applies Huffman coding. The results are reported for the predictor function SV7 which resulted in the best average compression ratio for the dataset.

- **JPEG-LS:** The standard of the JPEG committee for lossless and near-lossless compression of still images (ISO/IEC, 1999). The standard describes low-complexity predictive image compression algorithm with entropy coding using modified Golomb-Rice family. The algorithm is based on the LOCO-I algorithm (Weinberger, Seroussi, & Sapiro, 1996). We used the University of British Columbia implementation (version 2.2, ftp://ftp.netbsd.org/pub/NetBSD/packages/distfiles/jpeg_ls_v2.2.tar.gz).
- **JPEG2000:** A more recent JPEG committee standard describing an algorithm based on wavelet transform image decomposition and arithmetic coding (ISO/IEC, 2002). Apart from lossy and lossless compressing and decompressing of whole images, it delivers many interesting features such as progressive transmission, region of interest coding, and so forth (Christopoulos, Skodras, & Ebrahimi, 2000). We used the JasPer implementation by Adams (version 1.700.0, http://www.ece.uvic.ca/~mdadams/jasper/).
- **CALIC (Context-based Adaptive Lossless Image Compression):** A relatively complex predictive image compression algorithm using arithmetic entropy coding, which because of its usually good compression ratios is commonly used as a reference for other image compression algorithms (Wu, 1997). We used the implementation by Yang. In contrast to the other algorithms, CALIC is designed for grayscale images only. We therefore apply CALIC to each of the individual channels separately.

Lossless JPEG, JPEG-LS, and JPEG2000 are covered by international ISO standards whereas TIFF represents an industry standard. All algorithms except CALIC are incorporated into the medical imaging DICOM (Mildenberger et al., 2002; National Electrical Manufacturers Association, 2004) standard.

RETINAL IMAGE DATASET

A large set of over 800 colour retinal images captured at various ophthalmology centres was used in our experiments. The set is large both with regards to the number of images as well as with regards to the actual sizes of individual images. All images were initially obtained in uncompressed 24-bit RGB format. Images contain between 1.4 and 3.5 millions pixels and hence require between 4 and 10 MB of storage space.

In order to verify whether there are certain image classes that are especially susceptible to compression (or especially hard to compress), we divided the dataset into categories according to the following criteria:

1. **Retinal region:** The whole set is divided into three groups: *nasal*, *posterior*, and *temporal*. Evaluating the compression performance for individual subgroups will highlight whether any algorithms work especially well on any of these categories.
2. **Image size:** Here the images fall mainly into two categories: those images with about 1.4 million pixels (*small*) and those with about 3.4 million pixels (*large*). Compression ratios and compression speed are sometimes dependent on the image size, evaluating the individual size categories will hence confirm whether any such dependency exists for retinal images.

The two criteria are independent of each other (i.e., there exist both *small* and *large* images of all three regions in the dataset) and are hence tested separately. Details on how the images are divided into the individual categories are given in Table 1.

EXPERIMENTAL RESULTS

Experimental results were obtained on a HP Proliant ML350G3 computer equipped with two Intel Xeon 3.06 GHz (512 KB cache memory) processors and Windows 2003 operating system. Single-threaded applications of algorithms used for comparisons were compiled using Intel C++ 8.1 compiler. To minimise effects of the system load and the input-output subsystem performance, the algorithms were run several times; the time of the first run was ignored while the collective time of other

Related Content