J. Vis. Commun. Image R. 31 (2015) 253–265

Contents lists available at ScienceDirect

J. Vis. Commun. Image R.

journal homepage: www.elsevier.com/locate/jvci

Beyond pure quality: Progressive modes, region of interest coding, and real time video decoding for PDE-based image compression ☆

Pascal Peter *, Christian Schmaltz, Nicolas Mach, Markus Mainberger, Joachim Weickert

Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science, Campus, E1.7, Saarland University, 66041 Saarbrücken, Germany

A R T I C L E   I N F O

Article history:
Received 12 February 2015
Accepted 27 June 2015
Available online 8 July 2015

Keywords:
Partial differential equation
Diffusion
EED
Progressive mode
Region of interest coding
Image compression
Real time
Video decoding

A B S T R A C T

Compared to transform-based image compression methods such as JPEG2000, approaches based on partial-differential equations (PDEs) are in a proof-of-concept stage. Nevertheless, R-EED, a codec employing edge-enhancing anisotropic diffusion (EED) and rectangular subdivision, can surpass JPEG2000 quality-wise. However, today's requirements for compression algorithms go beyond pure compression performance. Codecs must also fulfil the feature requirements of specific applications such as online media or medical imaging. We propose three such features for the R-EED codec. By reordering grey values and exploiting the subdivision scheme, we incorporate a progressive mode into R-EED that can outperform JPEG and JPEG2000. Additionally, we show that rectangular subdivision is well-suited for region of interest coding and adapt the quality of image parts according to their importance. Finally, we propose a real-time video player that demonstrates how R-EED-based decoding can be performed efficiently. All of these extensions are compatible with each other and can be used simultaneously.

1. Introduction

Lossy image compression aims at reducing the file size of an image while degrading the image quality as little as possible. One of the most prominent and widely used lossy image compression algorithms today is JPEG [1], which is based on the discrete cosine transform. Its successor JPEG2000 [2], which uses Cohen–Daubechies–Feauveau wavelets, performs noticeably better, but is not yet widely supported.

Recently, these traditional transformation-based approaches have been challenged by novel compression methods that rely on an entirely different concept: Algorithms that use partial differential equations (*PDEs*) only store a subset of all image points and reconstruct the remainder of the image with PDE-based interpolation.

While there is a long history on research on feature-based image representations where homogeneous diffusion fills in missing information (see e.g. [3–7]), most of the early works do not consider applications in image compression. Only recently it has been shown that edge- or segment-based homogeneous diffusion approaches can beat JPEG2000 for cartoon-like images [8], as well as for depth maps [9–11]. Apart from these linear, feature-based approaches, there are some attempts to use nonlinear PDEs for image compression. Chan and Zhou [12] propose a variational approach with total variation regularisation to minimise oscillations in wavelet decompositions. Work by Solé et al. [13] evaluates different PDEs for compression of digital elevation maps. Liu et al. integrate inpainting into existing codecs [14]. The image compression algorithm of Galić et al. [15] uses points located on an adaptive triangulation and stores these locations as a binary tree. This method reconstructs the remainder of the image with edge-enhancing anisotropic diffusion (EED) [16]. Even with the improvements introduced in [17], the performance of this codec only lies between that of JPEG and JPEG2000 for medium to high compression ratios. The R-EED approach of Schmaltz et al. [18] builds upon these ideas. By introducing a number of novel concepts, e.g. using rectangular subdivisions instead of triangular ones, the obtained image compression codec yields better results. They can even surpass those of JPEG2000 for most compression ratios in images that are not dominated by texture. Schmaltz et al. demonstrated in [19] how one can further improve the compression quality of the R-EED algorithm.

However, a good reconstruction at high compression ratios alone is insufficient, since there are more requirements for an image compression algorithm in practice. Depending on the field in which the codec is applied, such requirements differ. For example in medical imaging, a high reconstruction quality in diagnostically relevant

☆ This paper has been recommended for acceptance by M.T. Sun.

* Corresponding author.

E-mail addresses: peter@mia.uni-saarland.de (P. Peter), schmaltz@mia.uni-saarland.de (C. Schmaltz), mach@mia.uni-saarland.de (N. Mach), mainberger@mia.uni-saarland.de (M. Mainberger), weickert@mia.uni-saarland.de (J. Weickert).

http://dx.doi.org/10.1016/j.jvcir.2015.06.017

regions is essential, while other regions can contain more errors. In web-based applications, codecs have to be able to deal with partially transferred files to reduce load times, and in time critical applications such as video playback, real-time capabilities are important. In this paper, we address such additional requirements for R-EED. In particular, we focus on progressive modes, region of interest coding, and real-time video decoding.

1.1. Related work

1.1.1. Progressive mode

The progressive mode, which is also referred to as embedded bitstreams or signal-to-noise ratio scalability, allows to generate a coarse preview from the beginning of a complete data stream only. This is especially advisable for applications in which only a limited bandwidth is available, e.g. when browsing a database of large images.

Progressive modes are readily available in most standard image compression codecs. Lossless image compression algorithms typically use interleaving or interlacing, which only changes the order in which pixels are stored. The well-known GIF file format [20], for example, divides the image into stripes with a height of eight pixels each, and stores the resulting lines in each stripe during multiple passes. PNG [21] employs a similar idea known as *Adam7*, which samples in both directions: It subdivides the image into 8×8 blocks and stores them in a total of seven passes.

The lossy mode of the JPEG standard provides several optional progressive modes: The progressive spectral selection reorders the transmission of the DC and AC components such that all low-frequent coefficients are transmitted first. Alternatively, one can successively approximate the stored coefficients by first storing the upper bits of the coefficients, while the lower bits follow later. It is possible, albeit unusual, to use both approaches at the same time. From a purely qualitative perspective, the so-called hierarchical mode provides typically the best results at low bit rates. This mode first stores a downsampled version of the image. The upsampled, bi-linearly interpolated version of this smaller image acts as a predictor for the next resolution. However, the total file size can increase up to one third with this mode [1].

In JPEG2000, images are always stored in progressive mode due to the *Embedded Block Coding with Optimal Truncation* (*EBCOT*) scheme. EBCOT encodes the bit-planes in three passes. In which pass a bit gets encoded depends on which coefficients have significant neighbours or are significant themselves.

So far, none of the PDE-based image compression methods discussed above include progressive modes. Note that standard progressive modes of existing image compression algorithms are not directly applicable to PDE-based methods, since data is only available at irregularly placed points.

1.1.2. Region of interest coding

There is no direct support for region of interest coding in JPEG. Nevertheless, a simple standard compliant idea proposed is to set additional DC coefficients to zero outside the region of interest [1].

JPEG2000 contains many methods for region of interest coding, namely a general scaling based approach, a bitplane-by-bitplane shift method (BbBShift), a max-shift method, a partial significant bitplane shift method (PSBShift) and a ROI coding through component priority (ROITCOP). Due to the large amount of approaches, we refer to [22,23] for details.

To the best of our knowledge, no PDE-based codec with ROI coding capabilities has been presented so far.

1.1.3. Real-time video decoding

Almost the whole existing body of work on PDE-based compression focuses on still image compression. However, there are some notable exceptions from the rule. Gao [24] used diffusion-based inpainting to compress optic flow fields for motion compensation in video coding. However, the actual compression of video data in this approach relies on block-coding with the discrete cosine transform. Therefore, the codec is still mainly transform-based. So far, real-time video decoding with fully PDE-based algorithms has only been implemented by Köstler et al. [25] and Baum [26]. Note that the work of Köstler et al. is based on linear diffusion which is computationally much less demanding than the high quality reconstructions of nonlinear anisotropic methods such as R-EED. Although the approach of Baum allows both linear and nonlinear anisotropic inpainting, it only achieves real-time decompression at the price of a noticeably degraded compression quality.

1.2. Goals

Consequently, the goal of this paper is threefold: First we give a detailed description how progressive modes can be incorporated into a current PDE-based image compression algorithm. This first part refines and expands work previously presented in [27]. Secondly, we extend our algorithm such that it includes a region of interest coding. This allows us to specify which parts of the image are important and should be saved with a higher precision. Finally, we introduce a PDE-based video decoder that allows real-time playback with nonlinear anisotropic diffusion inpainting.

Since R-EED is currently the most advanced PDE-based general purpose codec that can surpass JPEG2000 (see [18]), we use it as the basis of our approach. However, we have modified R-EED slightly to further enhance the performance of the algorithm independently of the novel features. Additionally, all three extensions are designed to be compatible, i.e. progressive mode and region of interest coding can both be applied simultaneously and are also applicable to video data.

1.3. Paper structure

First, we provide a brief overview on PDE-based compression and the baseline codec R-EED in Section 2. Section 3 contains a detailed discussion of the progressive mode for R-EED as well as comparisons to JPEG and JPEG2000. We introduce the second extension of R-EED, region of interest coding, in Section 4 and demonstrate its capabilities experimentally. In Section 5, we propose a real-time video-player based on R-EED compression and compare its performance to reference results obtained from classical solvers for PDE-based inpainting. Finally, we conclude the paper with a summary and an outlook in Section 6.

2. The baseline image compression codec

The R-EED codec [19] only stores the image values at a few selected locations, and employs PDE-based interpolation to reconstruct the missing data when loading the image. Thus, we first explain the ideas behind inpainting with PDEs. A description of the complete R-EED algorithm follows in Section 2.2.

2.1. PDE-based interpolation

Let $f : \Omega \to \mathbb{R}$ denote a grey value image with a rectangular image domain $\Omega \subset \mathbb{R}^2$. In PDE-based image interpolation, the grey values are only known at the locations of the *interpolation mask* $K \subset \Omega$ of the complete image f. The additional assumption that the pixels in $\Omega \setminus K$ are regular in some sense allows to reconstruct the unknown part of the image from the known grey values in K. For example, one might assume that the reconstructed image u is