Three-Dimensional Data Compression with Anisotropic Diffusion

Pascal Peter*

Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science,
Campus E1.7, Saarland University, 66041 Saarbrücken, Germany.
peter@mia.uni-saarland.de

Abstract. In 2-D image compression, recent approaches based on image inpainting with edge-enhancing anisotropic diffusion (EED) rival the transform-based quasi-standards JPEG and JPEG 2000 and are even able to surpass it. In this paper, we extend successful concepts from these 2-D methods to the 3-D setting, thereby establishing the first PDE-based 3-D image compression algorithm. This codec uses a cuboidal subdivision strategy to select and efficiently store a small set of sparse image data and reconstructs missing image parts with EED-based inpainting. An evaluation on real-world medical data substantiates the superior performance of this new algorithm in comparison to 2-D inpainting methods and the compression standard DICOM for medical data.

Keywords: 3-D image compression, edge-enhancing anisotropic diffusion, cuboidal subdivision, medical images

1 Introduction

Even with today's advances in storage capacity and transfer bandwiths, image compression remains an important area of research. In addition to 2-D images, great amounts of volumetric 3-D data is produced from a wide variety of recording sources such as CT or MRI scans. With JPEG 2000 [9] and DICOM [7] for medical imaging, both the 2-D and 3-D setting are dominated by transform-based approaches. However, a new family of image compression algorithms based on partial differential equations (PDE) has recently emerged [1, 6, 8]. Those methods have successfully challenged and in some cases surpassed the quality of the established codecs. In this paper, we aim to achieve three goals: 1. Introduce the first PDE-based compression codec explicitly designed for 3-D data. 2. Assess its potential in relation to established codecs. 3. Analyse the influence of 3-D diffusion on compression quality.

Related Work. PDE-based approaches [1, 6, 8] rely on the common idea to store only a small, systematically chosen subset of the image efficiently and reconstruct the missing parts by PDE-based interpolation. Methods based on edge-enhancing diffusion (EED) [1, 8] restrict the selection of known data for the sake of storage efficiency and compensate this with the powerful interpolation capacities of EED. Compression using edge information and homogeneous diffusion [6] is mostly efficient on cartoon-like

* recommended for submission to YRF2012 by Prof. Joachim Weickert

images. Our 3-D compression algorithm is modelled after R-EED [8], an algorithm that chooses sparse data with a rectangular subdivision scheme and fills in missing image parts with EED-based interpolation. To our best knowledge, PDE-based codecs have not been used for spatial 3-D data before our approach. However, there are video codecs that rely on inpainting techniques [2, 5, 4, 11]. In particular, PDE-based video compression approaches based on [1] were proposed in [2] and [4].

Our Contribution. We extend the *rectangular* 2-D subdivision scheme R-EED [8] to 3-D, thereby proposing the first PDE-based 3-D compression codec. We introduce *cuboidal* subdivision as a natural 3-D extension of R-EED's rectangular subdivision scheme in order to store a sparse subset of the 3-D data. The rest of the image is reconstructed using 3-D EED inpainting. Following R-EED's naming scheme, the new codec is called C-EED. Furthermore, we analyse the influence of the 3-D diffusion step on compression quality by identifying all sources for performance gains of C-EED in comparison to R-EED applied to a sequence of 2-D image slices. Finally, we compare the overall performance of C-EED, its predecessor R-EED and the DICOM standard [7] on medical data.

2 Three-Dimensional EED-Based Inpainting

PDE-based inpainting acts as the foundation of our compression algorithm. In general, inpainting denotes the reconstruction of missing image parts from known data. We exploit this concept for lossy compression by keeping only fractions of the original data and filling in the rest via inpainting. The following sections cover both a continuous model for inpainting as well as the discretisation that we apply in our method.

Continuous Inpainting Model. In the following, we consider 3-D grey value images as sufficiently smooth functions that map image coordinates from the cuboidal image domain $\Omega \subset \mathbb{R}^3$ to a continuous grey value range. The goal of the inpainting process is to use a set of known data, the so-called inpainting mask $M \subset \Omega$, to reconstruct the contents of the inpainting domain $\Omega \setminus M$. Based on its superior performance in 2-D [8] we choose edge-enhancing anisotropic diffusion (EED) [10] as the basis for the inpainting algorithm. EED inpainting is defined by a boundary value problem:

$$\begin{aligned} \partial_t u = \operatorname{div}(\boldsymbol{D}\boldsymbol{\nabla} u) \quad &\text{on} \quad \Omega \setminus M \times (0,\infty), \\ u(\boldsymbol{x},t) = f(\boldsymbol{x}) \quad &\text{on} \quad \partial M \times (0,\infty), \\ \langle \boldsymbol{D}\boldsymbol{\nabla} u, \boldsymbol{n}\rangle = 0 \quad &\text{on} \quad \partial\Omega \times (0,\infty). \end{aligned} \tag{1}$$

This system describes the evolution of the inpainted image u under EED with initial data f and mixed boundary conditions: on the image boundaries $\partial\Omega$, Neumann conditions are applied, while the boundary ∂M of the inpainting mask is treated with Dirichlet conditions. Here, $\boldsymbol{\nabla} u$ denotes the spatial gradient and div the spatial divergence operator.

The diffusion tensor $\boldsymbol{D} \in \mathbb{R}^{3\times 3}$ defines the characteristic properties of EED: Diffusion is performed only along image edges, not across them. This behaviour is achieved by carefully designing $\boldsymbol{D}$ in terms of its eigenvectors $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3$ and eigenvalues μ_1, μ_2, μ_3. Each eigenvalue determines the amount of diffusion performed in the direction of the corresponding eigenvector. In 2-D, $\boldsymbol{\nabla} u$ is perpendicular to the local edge direction.

Thus, the relevant directions are defined parallel to $\boldsymbol{\nabla} u$ (across the edge) and parallel to $\boldsymbol{\nabla} u$ (edge direction). In 3-D, $\boldsymbol{\nabla} u$ still defines the direction across the edge. However, diffusion along the edge is performed in the 2-D plane with normal vector $\boldsymbol{\nabla} u$.

$$\boldsymbol{v}_1 \parallel \boldsymbol{\nabla} u_\sigma \;\text{ with }\; \mu_1(\boldsymbol{\nabla} u_\sigma) := \frac{1}{\sqrt{1 + \frac{|\boldsymbol{\nabla} u_\sigma|^2}{\lambda^2}}} \tag{2}$$

$$\boldsymbol{v}_2 \perp \boldsymbol{\nabla} u_\sigma, \boldsymbol{v}_3 \perp \boldsymbol{\nabla} u_\sigma, \boldsymbol{v}_2 \perp \boldsymbol{v}_3 \;\text{ with }\; \mu_2 = \mu_3 := 1 \tag{3}$$

Equation (3) ensures unrestricted diffusion along the edges. As in [8], diffusion across edges is penalised proportional to the edge detector $|\boldsymbol{\nabla} u_\sigma|^2$ by the Charbonnier diffusivity in Equation (2). The free parameter λ can be used to adjust the edge preservation of EED and thereby the contrast of the diffusion results. The presmoothing of the image u by convolution $u_\sigma := K_\sigma * u$ with a Gaussian kernel K_σ with standard deviation σ robustifies the edge detection.

Discretisation. We discretise the system (1) by a finite difference approximation on a spatiotemporal grid with grid sizes h_x, h_y, h_z and h_t. Solving the resulting equations for the evolved image after one discrete time step yields a so-called explicit scheme which allows to implement EED with a fast explicit diffusion approach [3]. In order to account for real-world images with nonuniform voxel size, we adapt the spatial grid sizes accordingly.

3 The C-EED Codec

Our C-EED algorithm is modeled after the 2-D codec R-EED [8]. Both methods share the common idea of reconstructing the image from a sparse subset of the image data. One of the most prominent changes is the 3-D inpainting described in the previous section. However, the inpainting process alone does not necessarily yield a good reconstruction. The choice of the inpainting mask M is an equally important problem.

Cuboidal Subdivision. Selecting the inpainting mask M is guided by three competing factors: 1. The choice of mask points influences the *quality* of the inpainting results. 2. *Storage efficiency* must be considered. This includes the storage of mask point locations and the influence on efficient grey value storage by entropy coding. 3. The *complexity* of finding the optimal mask for a given compression ratio is large. Paying tribute to all of these factors, the choice of mask points is restricted to an adaptive grid that allows to encode mask point locations in a binary tree. This tree is constructed by a threshold-based subdivision scheme. First, a static point pattern is chosen, in our case the corners and center of the cuboidal image domain. These points are added to the inpainting mask M and used to compute a preliminary reconstruction result. If the reconstruction error exceeds an error threshold E, the current image is split in the middle of its largest side into two cuboidal subimages of equal size. Mask points are added in these subimages according to the point pattern, and the splitting process is repeated for each subimage until all subimages respect the threshold E. Each split is represented by one node of the corresponding binary subdivision tree. The error threshold is adapted to the tree depth by the formula $E = a\ell^d$, where d is the tree depth and a and ℓ are free parameters. This subdivision scheme is illustrated by Fig. 1.

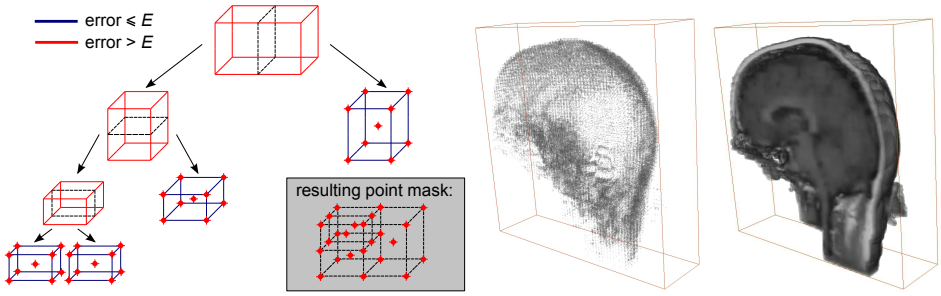

Fig. 1. Cuboidal Subdivision. The schematic representation on the left shows the construction of an inpainting mask with cuboidal subdivision. The two medical volume images are examples for an inpainting mask (middle) and a C-EED reconstruction (right).

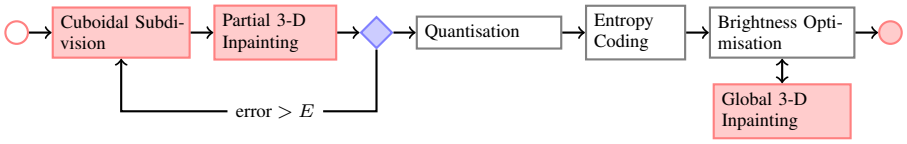

Fig. 2. Encoding with C-EED. The flow chart above describes the C-EED steps performed for image encoding. Highlighted steps mark prominent differences between C-EED and R-EED.

Encoding. The additional encoding steps that were not discussed in previous sections are natural extensions of the compressions steps of R-EED [8]. First, a *quantisation* of the grey value domain is performed in order to increase the efficiency of the final entropy coding step. Afterwards, the inpainting mask is computed by *cuboidal subdivision* with inpainting operations that are restricted to subimages. The inpainting quality can be increased further by *brightness optimisations*: grey values of the inpainting mask are modified in order to improve inpainting results. While this introduces an error to the mask points, the quality gain in the inpainting domain significantly outweighs this loss. Finally, the data is stored to the hard disk. The file header contains the image size and the compression parameters EED-contrast λ, number of quantised grey values q and swapping parameter s (see decoding step). Additionally the subdivision tree is saved in binary code. The grey value data is stored losslessly by *entropy-coding*. Note that the encoding procedure in Fig. 2 assumes that all free parameters are given. In practice, all parameters are optimised to achieve results of the highest quality possible.

Decoding. The compression parameters, subdivision tree and grey values are decoded. From these values, the inpainting mask is constructed and with a single inpainting step, the decompressed image is obtained. As a post processing step, an *interpolation swap* can be performed: the role of known and unknown data is inverted and thereby mask point degradation by quantisation and brightness optimisation can be attenuated. The new inpainting domain consists of spheres with radius s around all mask points.

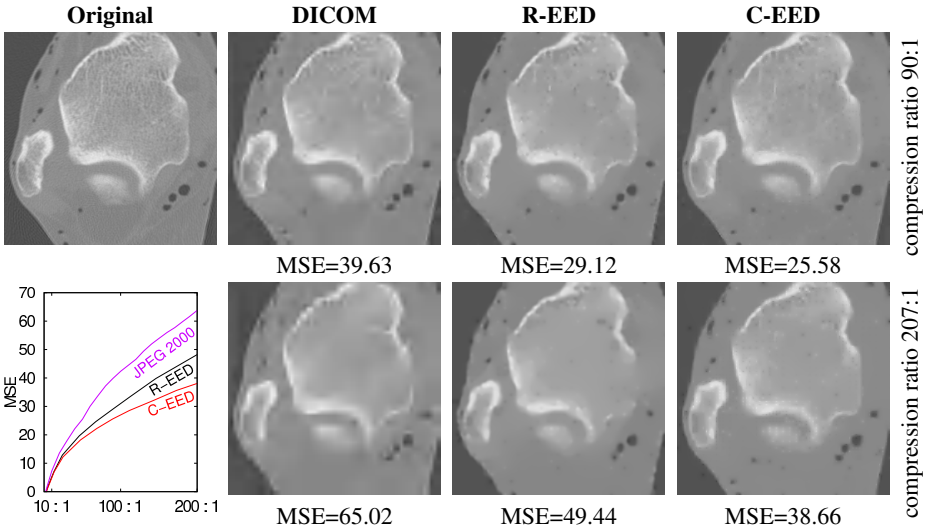

Fig. 3. Images obtained with DICOM, R-EED and C-EED with compression rates close to 90:1 (top) and 207:1 (bottom). The test image (size $256 \times 256 \times 64$) consists of the first 64 slices of a trabular bone CT provided by Wiro Niessen. The images above display the last 2-D slice of the test image. Smaller MSE values are better.

4 Experiments

Experiment Design. Due to space constraints we present a small representative subset of our experiments with C-EED, R-EED and DICOM [7] on spatial 3-D data. Image quality is assessed with respect to the mean squared error (MSE). For compression with DICOM the most efficient method in the standard, JPEG 2000, is used. The application of 2-D codecs to 3-D data is usually achieved by separation of 3-D data into 2-D slices. On of our main goals is to assess the influence of 3-D EED-based inpainting in comparison to slice-wise inpainting with 2-D EED. In order to isolate the 3-D diffusion's contribution, further sources for efficiency improvements must be incorporated into R-EED: 1. *Header redundancies* appear due to duplicate header data in each compressed 2-D slice. 2. *Global entropy coding* is, in most cases, more efficient than individual coding for each slice. Therefore, all experiments were conducted with a modified version of R-EED that contains all improvements of C-EED except for 3-D EED.

Results. The results in Fig. 3 reveal that quality gaps between the different approaches in terms of the MSE are amplified for increasing compression rates. For medium to high compression rates the quality of DICOM degrades much faster than the PDE-based approaches. In the same range C-EED shows significant advantages over R-EED. Visually, DICOM suffers from heavy block artifacts, while R-EED and C-EED produce smooth results that only lose detail with increasing compression rate. Similar results were observed on all medical test data. On synthetic data with high redundancy in one dimension, C-EED yields improvements over R-EED of up to 50%.

Overall, the experiments demonstrate the application of 3-D inpainting in C-EED is very well suited for the compression of spatial 3-D data and offers significant advantages over R-EED and DICOM.

5 Conclusion

In this paper we introduced the first PDE-based 3-D image compression algorithm and demonstrated its superior performance on 3-D spatial data in comparison to DICOM and R-EED. In particular, we analysed the influence of 3-D diffusion on image quality and verified that it offers significant advantages over 3-D compression based on 2-D diffusion.

In future work, the potential of 3-D diffusion on other 3-D data should be closely investigated. In particular, C-EED should be well suited for seismic data sets with huge file sizes (1 TB and more) where the aim is very high compression rates. Furthermore, the potential of C-EED compression on spatiotemporal data should be investigated, since this could pave the way for competitive PDE-based video codecs.

References

1. Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., Seidel, H.P.: Image compression with anisotropic diffusion. Journal of Mathematical Imaging and Vision 31(2–3), 255–269 (2008)
2. Gao, Q.: Low bit rate video compression using inpainting PDEs and optic flow. Master's thesis, Dept. of Mathematics and Computer Science, Saarland University, Germany (2008)
3. Grewenig, S., Weickert, J., Bruhn, A.: From box filtering to fast explicit diffusion. In: Goesele, M., Roth, S., Kuijper, A., Schiele, B., Schindler, K. (eds.) Pattern Recognition, Lecture Notes in Computer Science, vol. 6376, pp. 533–542. Springer, Berlin (2010)
4. Köstler, H., Stürmer, M., Freundl, C., Rüde, U.: PDE based video compression in real time. Tech. Rep. 07-11, Lehrstuhl für Informatik 10, Univ. Erlangen–Nürnberg, Germany (2007)
5. Liu, D., Sun, X., Wu, F., Li, S., Zhang, Y.Q.: Image compression with edge-based inpainting. IEEE Transactions on Circuits, Systems and Video Technology 17(10), 1273–1286 (2007)
6. Mainberger, M., Weickert, J.: Edge-based image compression with homogeneous diffusion. In: Jiang, X., Petkov, N. (eds.) Computer Analysis of Images and Patterns, Lecture Notes in Computer Science, vol. 5702, pp. 476–483. Springer, Berlin (2009)
7. National Electrical Manufacturers Association: Digital Imaging and Communications in Medicine (DICOM) – Part 5 Data Structures and Encoding. PS 3.5-2004 (2004)
8. Schmaltz, C., Weickert, J., Bruhn, A.: Beating the quality of JPEG 2000 with anisotropic diffusion. In: Denzler, J., Notni, G., Süße, H. (eds.) Pattern Recognition, Lecture Notes in Computer Science, vol. 5748, pp. 452–461. Springer, Berlin (2009)
9. Taubman, D.S., Marcellin, M.W. (eds.): JPEG 2000: Image Compression Fundamentals, Standards and Practice. Kluwer, Boston (2002)
10. Weickert, J.: Scale-space properties of nonlinear diffusion filtering with a diffusion tensor. Tech. Rep. 110, Laboratory of Technomathematics, University of Kaiserslautern, Germany (1994)
11. Zhu, C., Sun, X., Wu, F., Li, H.: Video coding with spatio-temporal texture synthesis and edge-based inpainting. In: Proc. 2008 IEEE International Conference on Multimedia and Expo. pp. 813–816. Hannover, Germany (2008)