Report from Dagstuhl Seminar 16462

Inpainting-Based Image Compression

Edited by

Christine Guillemot[1], Gerlind Plonka-Hoch[2], Thomas Pock[3], and Joachim Weickert[4]

1 **INRIA – Rennes, FR,** christine.guillemot@inria.fr
2 **Universität Göttingen, DE,** plonka@math.uni-goettingen.de
3 **TU Graz, AT,** pock@icg.tugraz.at
4 **Universität des Saarlandes, DE,** weickert@mia.uni-saarland.de

Abstract

Inpainting-based image compression is an emerging paradigm for compressing visual data in a completely different way than popular transform-based methods such as JPEG. The underlying idea sounds very simple: One stores only a small, carefully selected subset of the data, which results in a substantial reduction of the file size. In the decoding phase, one interpolates the missing data by means of a suitable inpainting process. It propagates information from the known data into the areas where nothing has been stored, e.g. by solving a partial differential equation or by clever copy-and-paste mechanisms.

Inpainting-based codecs (coders and decoders) are more intuitive than transform-based ones, they are closer to biological mechanisms in our brain, and first results show that they may offer promising performance for high compression rates. However, before these ideas become practically viable, a number of difficult fundamental problems must be solved first. They involve e.g. the selection of the data and the inpainting operator, coding strategies, and the search for highly efficient numerical algorithms. This requires a collaborative effort of experts in data compression, inpainting, optimisation, approximation theory, numerical algorithms, and biological vision. In this Dagstuhl seminar we have brought together leading researcher from all these fields for the first time. It enabled a very fruitful and inspiring interaction which will form the basis for future progress.

Seminar November 13–18, 2016 – http://www.dagstuhl.de/16462

1998 ACM Subject Classification E.4 Coding and Information Theory, F.2.1 Numerical Algorithms and Problems, G.1.1 Interpolation, G.1.2 Approximation, G.1.6 Optimization, G.1.8 Partial Differential Equations, G.2 Discrete Mathematics, I.4.2 Compression (Coding)

Keywords and phrases approximation, inpainting, interpolation, lossy image compression, optimisation, partial differential equations (PDEs), sparsity

Digital Object Identifier 10.4230/DagRep.6.11.90

Edited in cooperation with Sarah Andris

1 Executive Summary

Joachim Weickert

Since the amount of visual data is rapidly increasing, there is a high demand for powerful methods for compressing digital images. A well-known example is the lossy JPEG standard that is based on the discrete cosine transform. Unfortunately its quality deteriorates substantially for high compression rates, such that better alternatives are needed.

The goal of this seminar was to pursue a completely different strategy than traditional, transform-based codecs (coders and decoders): We studied approaches that rely on so-called inpainting methods. They store only a small, carefully selected subset of the image data. In the decoding phase, the missing data is reconstructed by interpolation with partial differential equations (PDEs) or by copying information from patches in other image regions. Such codecs allow a very intuitive interpretation, and first experiments show their advantages for high compression rates where they can beat even advanced transform-based methods.

However, inpainting-based codecs are still in an early stage and require to solve a number of challenging fundamental problems, in particular:

1. Which data gives the best reconstruction?
2. What are the optimal inpainting operators?
3. How should the selected data be encoded and decoded?
4. What are the most efficient algorithms for real-time applications?

These problems are highly interrelated. Moreover, they require interdisciplinary expertise from various fields such as image inpainting, data compression and coding, approximation theory, and optimisation. To design these codecs in an optimal way, one must also understand their connections to related areas such as sparsity and compressed sensing, harmonic analysis, scattered data approximation with radial basis functions, and subdivision strategies.

Our seminar constituted the first symposium on this topic. It brought together 29 researchers from 11 countries, covering a broad range of expertise in the different fields mentioned above. Many of them have met for the first time, which resulted in a very fruitful interaction.

In order to have a good basis for joint discussions, first all participants introduced themselves and briefly described their background and interests. Then the seminar proceeded with six tutorial talks (45 minutes plus 15 minutes discussion), given by the four organisers as well as by Simon Masnou and Nira Dyn. In this way all participants could acquire a general overview on the achievements and challenges of inpainting-based image compression and its various aspects such as coding, inpainting, convex optimisation, subdivision, and computational harmonic analysis.

Afterwards we decided to cluster the talks thematically into six sessions, each consisting of 3–4 talks (ca. 30 minutes plus 15 minutes discussion) and lasting half a day:

1. Harmonic Analysis
   (talks by Gerlind Plonka-Hoch, Naoki Saito, and Hao-Min Zhou)
2. Approximation Theory
   (talks by Martin Buhmann, Armin Iske, Nira Dyn, and Tomas Sauer)
3. Inpainting
   (talks by Aurelien Bourquard, Carola-Bibiane Schönlieb, and Yann Gousseau)
4. Compression
   (talks by Gene Cheung, Joan Serra Sagrista, and Claire Mantel)
5. Optimisation of Data and Operators
   (talks by Zakaria Belhachmi, Laurent Hoeltgen, Peter Ochs, and Pascal Peter)
6. Algorithms, Biological Vision, and Benchmarking
   (talks by Jalal Fadili, Johannes Ballé, and Sarah Andris)

These sessions triggered interesting discussions during the talks, in the breaks, and in the evening, and they allowed the different communities to learn many new things from each other.

Our program featured also an evening panel discussion on open research questions on the interface between image inpainting and image compression. It was a lively interaction between the five panel members and the audience, involving also controversial statements and views about the future of inpainting-based codecs.

The participants had a very positive impression of this seminar as an inspiring forum to bring together different fields. As a consequence, this symposium also created several new collaborations, e.g. regarding interpolation with radial basis functions, subdivision-based coding, and diffusion-based coding. There was a general consensus that it would be desirable to have another seminar on this topic in 2–3 years. Moreover, it is planned to compile a related monograph which will be the first in its field.

2 Table of Contents

3 Overview of Talks

3.1 A Benchmark for Inpainting-Based Image Reconstruction and Compression

Sarah Andris (Universität des Saarlandes, DE)

Joint work of Sarah Andris, Leif Bergerhoff, Joachim Weickert

Well-designed benchmarks have been able to push research forward in various fields by providing a basis for fair comparison. To the best of our knowledge, there is no benchmark for inpainting-based image compression that provides suitable sets of images and allows to compare and rank different methods. Our goal is to introduce such a benchmark. To this end, we present a webpage that contains preliminary test sets and allows users to test the ability of masks and corresponding grey or colour values to reconstruct images from these sets. Rather than introducing a complete benchmark, this talk intends to establish a foundation for discussion on the design of this webpage, offering different possibilities and extensions. We are convinced that an acknowledged benchmark in this area is able to foster fair comparisons and can promote reproducible research.

3.2 Nonlinear Image Representations with Cascaded Local Gain Control

Johannes Ballé (New York University, US)

Joint work of Johannes Ballé, Valero Laparra, Eero Simoncelli

Local gain control is found throughout biological sensory systems. Modeled as an operation known as divisive normalisation, it can be implemented as an invertible nonlinear transform, and has several interesting properties useful to image processing. It can be used to Gaussianise image densities, eliminating statistical dependencies and providing probabilistic image models that outperform sparse representations. It can mimic human perception of visual distortions, often manifesting as masking effects, and additionally can be used for image compression. The image degradations due to compression appear reminiscent of the results of edge-preserving smoothing methods, such as anisotropic diffusion.

3.3 What is a Right Operator in Inpainting?

Zakaria Belhachmi (University of Mulhouse, FR)

The choice of a "right" model/operator for the reconstruction in inpainting is still a challenging problem. We present some ideas on modeling with diffusion operators for geometry inpainting. As framework we choose variational methods based on partial differential equations (PDEs) and their discrete counterparts. The objective is to provide a closed loop continuous to discrete model. The loop consists of an initial family of simple PDEs depending on some parameters – "corrected" at the discrete level – which yields (in the Gamma-convergence sense) a sophisticated limit model that captures the jump set of the restored image. We also discuss inpainting-based compression within this approach.

3.4 Variations on Edge-Enhancing Diffusion: Sparse Interpolation and Beyond

Aurélien Bourquard (MIT – Cambridge, US)

In this talk, we shall introduce an approach to interpolate images from sparse sets of samples. The proposed method aims at preserving image edges and is thus inspired from edge-enhancing diffusion (EED). Meanwhile, one important characteristic of our algorithm is that it follows variational principles. Specifically, it involves the minimisation of successive quadratic-cost functionals, which relates to iteratively-reweighted-least-squares (IRLS) methods. Considering several image examples, we shall illustrate how such an approach is able to produce high-quality results while maintaining computational efficiency, even when as little as 2% of the image samples are available. Finally, we shall discuss how one could leverage the enhanced modularity of the proposed algorithm's structure for adaptations to wider classes of problems, including compressed sensing.

3.5 Convergence Properties of Multiquadrics with Parameters

Martin Buhmann (Universität Gießen, DE)

Radial basis function methods are a useful approach especially to multivariable approximation, e.g., by interpolation, but also using the successful alternative called quasi-interpolation. This is all based on the idea that one can employ shifts of a single radially symmetric function to form spaces of approximating functions in any number of variables. This univariate function can be a decaying function such as an exponential, but unbounded radial basis functions such as multiquadrics are successful, too.

The initial approach to choosing which approximants to take, given possibly a very large number of scattered data in high dimensions, is usually by interpolation. Therefore one uses shifts of a radial basis function, such as the multiquadrics or Gauss- or Poisson-kernel, all of which are related to solving certain partial differential equations and to minimising some particular semi-norms, to interpolate the mentioned data. This of course raises the question of solvability of the problem and of the efficient computability.

We will discuss in this talk mainly multiquadric interpolation, its convergence properties when centres become dense in multiple dimensions, and their behaviour when parameters of the multiquadric function change.

3.6 Graph Signal Processing for Image Coding and Restoration

Gene Cheung (National Institute of Informatics – Tokyo, JP)

Graph signal processing (GSP) is the study of discrete signals that live on structured data kernels described by graphs. By allowing a more flexible graphical description of the underlying data kernel, GSP can be viewed as a generalisation of traditional signal processing techniques that target signals in regular kernels – e.g., an audio clip sampled periodically in time – while still providing a frequency domain interpretation of the observed signals. Though an image is a regularly sampled signal on a 2-D grid, one can nonetheless consider an image or an image patch as a graph-signal on a sparsely connected graph defined signal-dependently. Recent GSP works have shown that such an approach can lead to a compact signal representation in the graph Fourier domain, resulting in noticeable gain in image compression and restoration. Specifically, in this talk I will overview recent advances in GSP as applied to image processing. We will first describe how a Graph Fourier Transform (GFT) – a generalisation of known transforms like Discrete Cosine Transform (DCT) and Asymmetric Discrete Sine Transform (ADST) – can be defined in a signal-dependent manner and leads to compression gain over traditional DCT for piecewise smooth images. We will then describe how suitable graph-signal smoothness priors can be constructed for a graph-based image denoising algorithm for piecewise smooth images. Similar graph-signal smoothness priors can also be used for other image restoration problems, such as de-quantisation of compressed JPEG images.

3.7 Linear and Nonlinear Subdivision Schemes – An Overview

Nira Dyn (Tel Aviv University, IL)

Subdivision schemes are processes for the generation of curves/surfaces from discrete data by repeated refinements. The implementation of these schemes is simple, but their mathematical analysis is rather involved. This talk will present briefly the "classical" case of linear univariate schemes refining points for generating curves. In particular, two basic schemes will be considered. A short introduction to bivariate schemes generating surfaces, together with a specific example will also be discussed.

The talk will conclude with an extension of the 4-point linear scheme to a nonlinear scheme, capable of handling non-uniform geometries, together with a few examples of its performance.

3.8 IPOL: Image Processing On Line

Gabriele Facciolo (ENS – Cachan, FR)

It is clear that software is important in modern science. It is used to solve real scientific problems in a variety of fields ranging from physics, to medicine, and economics. However, research software is not like general purpose equipment (like telescopes or compilers). Research software is usually made by scientists for scientists, tailored to a concrete experimental process. So this code is not meant to be released nor published and even less reviewed. Nevertheless, the research code is very important, since it embodies the process itself. So the research software is key for attaining a truly reproducible research.

In this talk I am going to share our experience with IPOL: an on-line journal for image processing, a journal that publishes algorithm descriptions along with the code that implement them. IPOL follows the open access and reproducible research models. Since it is a journal, the code is also peer-reviewed to verify that it corresponds to the algorithm description.

The distinguishing characteristic of IPOL articles is that each article also includes an on-line demo that permits to easily test the code with new data and parameters from the web, and a public archive that stores all original test data used in the demo. Both the demo and the archive provide great insight into the algorithm inner workings and limitations, facilitating the experimentation and leading in the end to a stricter verification of the algorithm.

3.9 Variational Analysis Meets Stochastic Differential Equations: Application to Inverse Problems

Jalal Fadili (Caen University, FR)

Joint work of Christophe Chesneau, Jalal Fadili, Tùng Luu

Unlike regularised estimators which amount to solving an optimisation problem, computing the posterior conditional mean estimator corresponds to an integration problem which becomes very involved to solve analytically or even numerically in high-dimension. An alternative is to approximate it via a Markov chain Monte-Carlo (MCMC) method which consists in sampling from the target distribution by constructing an appropriate Markov chain, and to compute sample path averages based on the output of the Markov chain. In this work, we study an exponentially weighted aggregator/estimator. To compute this estimator, we propose a family of forward-backward-type proximal Langevin Monte-Carlo algorithms to efficiently sample from the target distribution (which is not smooth nor log-concave) and derive its guarantees. We illustrate it on various inverse problems including inpainting.

3.10 Patch-Based Video Inpainting

Yann Gousseau (Telecom ParisTech, FR)

Joint work of A. Almansa, M. Fradet, Y. Gousseau, T. T. Le, S. Masnou, A. Newson, P. Pérez

Video inpainting is a challenging task for several reasons, including: our sensibility to temporal discontinuity or fluctuations, the volume of data involved, the presence of complex motions and moving textures, etc. To the best of our knowledge, the most efficient approaches to this problem rely on heuristics to minimise a global, patch-based energy, as initially introduced by Y. Wexler et al. about ten years ago.

In this talk, I will first review these approaches, before presenting some recent contributions that can yield efficient reconstructions for complex and high resolution videos. Eventually I will show failure examples and open problems.

3.11 Overview of Image Compression: Basics and Research Issues

Christine Guillemot (INRIA – Rennes, FR)

Image and video compression is now an old research topic, spanning three decades from the mid 80s. Research in the field has produced an innumerable amount of algorithms and techniques. This talk will start with a few source coding basics, and we will see how these basics have guided the design of key coding operators. Focusing on two operators, transforms and prediction, the talk will then present advances from isotropic to anisotropic transforms, to sparse and graph-based transforms. We will then move to the problem of prediction which is analogous to inpainting, and present a prediction approach based on sparse models. We will finish by presenting a novel coding architecture building upon concepts of epitomes and based on inpainting.

3.12 Optimising Data for PDE-Based Inpainting and Compression

Laurent Hoeltgen (BTU Cottbus, DE)

Finding good reconstruction data is a key problem for PDE-based inpainting within the context of lossy image compression. Not only the location of important pixels but also their corresponding colour values should be optimal. In this talk we discuss how the spatial and tonal optimisation are related and provide strategies to retrieve this data. Finally, we will discuss additional optimisation challenges that occur during the encoding of the data.

3.13 Geometrical Methods for Adaptive Approximation of Image and Video Data

Armin Iske (Universität Hamburg, DE)

Joint work of Laurent Demaret, Nira Dyn, Armin Iske, Wahid Khachabi

We survey more recent and less recent results concerning adaptive approximation of images and videos. The utilised approximation methods rely on linear splines over anisotropic simplices and adaptive thinning, a greedy point removal scheme for multivariate scattered data. We will address both computational and theoretical aspects of adaptive thinning, including optimal N-term approximations.

3.14 Compression of Infrared Images

Claire Mantel (Technical University of Denmark – Lyngby, DK)

Advances in the manufacturing of thermal sensors (LWIR) have led to increased accessibility, improved precision, and higher resolution. IR images differ from visible light ones by their characteristics (due to the formation process of IR images) but also by their bitdepth (most IR sensors are 14b). This implies either tone mapping the images to 8b or high bitdepth compression. We will present a preliminary study on the compression of IR images.

3.15 Inpainting: An Overview

Simon Masnou (University Claude Bernard – Lyon, FR)

The purpose of this overview talk is to review various approaches which have been proposed to address the inpainting problem: geometric, patch-based, variational, or harmonic methods, and their combinations.

Geometric approaches consist in advecting or diffusing the information from outside to inside the inpainting domain, using variational or PDE models. Such methods are usually efficient for reconstructing geometric informations, but are not adapted to texture reconstruction. On the contrary, patch-based methods were initially designed for texture synthesis, and have been later adapted to image inpainting. Roughly speaking, they are based on copying and pasting (or averaging) similar "patches", i.e. pixel neighbourhoods. Modern approaches combine both geometric and patch-based principles, and rely crucially on a multiscale strategy. Excellent results can be obtained, although stability cannot be guaranteed: slightly different inpainting domains or parameters may yield very different results. A third category of methods performs very well for sparse inpainting domains: harmonic approaches, which use image decomposition on various dictionaries, for instance one adapted to geometry, another one adapted to texture, etc.

Numerous results for 2-D or 3-D images will be shown during the talk. The challenges of video inpainting will also be illustrated.

3.16 Optimising Data for Anisotropic Diffusion Inpainting with an Optimal Control Approach

Peter Ochs (Universität Freiburg, DE)

A key problem for inpainting processes based on partial differential equations (PDEs) in image compression applications is to find optimal locations and values of the stored grey values. Even for homogeneous diffusion models this is a challenging task. An optimal control approach is one of the best strategies to find optimal locations. Advanced anisotropic inpainting operators such as edge-enhancing diffusion (EED) yield a high reconstruction quality. This talk discusses an optimal control approach with the EED inpainting operator.

3.17 Gradients versus Grey Values for Sparse Image Reconstruction and Inpainting-Based Compression

Pascal Peter (Universität des Saarlandes, DE)

Joint work of Sebastian Hoffmann, Enric Meinhardt-Llopis, Pascal Peter, Markus Schneider, Joachim Weickert

Inpainting-based compression codecs store images in terms of a small fraction of pixels and interpolate the missing image parts for decoding. Recently, Brinkmann et al. (2015) have suggested an alternative sparse representation: They propose to store gradient data instead of grey values. However, the potential of this idea for PDE-based reconstruction and compression still needs to be evaluated.

The goal of this talk is to provide a fair comparison of gradient and grey value data for homogeneous diffusion inpainting. To this end, we establish a framework for optimising and encoding the known data. It allows a fair comparison of the grey value and the gradient approach, as well as combinations of both kinds of known data. In this talk, we distinguish between two separate assessments of quality: On one hand, pure reconstruction quality can be analysed by inpainting images from comparable amounts of known data for each method. On the other hand, compression quality is obtained by also taking storage costs into account by specifying a full codec for each approach.

Our evaluation shows that gradient-based reconstructions can improve the visual quality in a pure reconstruction setting: They avoid singularities involved in the reconstructions from grey values. Surprisingly, this advantage does not carry over to compression due to high sensitivity of gradient data to quantisation. Combining both kinds of data can attenuate this drawback of the gradients, but overall, pure grey-value reconstructions still remain the better choice for compression. This reveals an important general principle for inpainting-based compression: Good reconstruction quality from sparse data only leads to good compression performance, if the approach is also robust w.r.t. lossy compression steps.

3.18 Computational Harmonic Analysis Tools for Image Compression and Inpainting

Gerlind Plonka-Hoch (Universität Göttingen, DE)

In this talk we will give an overview on basic ideas for image compression and inpainting using transforms from computational harmonic analysis.

Standard compression schemes like JPEG are essentially based on data decorellation using the cosine transform or tensor-product wavelet transforms. In recent years, many further adaptive and non-adaptive transforms have been proposed that can be used for sparse data representation and image compression.

The talk shows some examples of adaptive wavelet transforms as well as non-adaptive directional wavelet frames. Adaptive transforms can be directly used for image compression. They need a careful design to ensure stability of the inverse transform or a precise analysis of occurring adaptivity costs. Nonadaptive redundant frames, as e.g. curvelets or shearlets, cannot be directly used for compression but may serve as a tool to construct suitable regularisation terms if we consider solving optimisation problems for inpainting.

3.19 Convex Optimisation and Inpainting: A Tutorial

Thomas Pock (TU Graz, AT)

In this talk, we will discuss recent trends in convex optimisation for image processing and in particular for image inpainting. First, we will discuss popular non-smooth convex models such as the total variation, higher-order extensions, and also curvature depending models. Then, we will see how these models can be applied for image inpainting problems and finally we will discuss efficient first-order algorithms to minimise these models.

3.20 Polyharmonic Local Cosine Transforms for Improving JPEG-Compressed Images

Naoki Saito (University of California – Davis, US)

I will discuss two image compression-decompression schemes that reproduce images with better visual fidelity, less blocking artifacts, and better PSNR, particularly in low bit rates, than those processed by the JPEG Baseline method at the same bit rates. We got the patents on these algorithms both in Japan (2009) and the US (2011). The additional computational cost is small, i.e., linearly proportional to the number of pixels in an input image. The first method, the "full mode" polyharmonic local cosine transform (PHLCT), modifies the encoder and decoder parts of the JPEG Baseline method. The goal of the full mode PHLCT is to reduce the code size in the encoding part and reduce the blocking artifacts in the decoder part. The second one, the "partial mode" PHLCT (or PPHLCT for short), modifies only the decoder part, and consequently, accepts JPEG files, yet decompresses them with higher quality with less blocking artifacts. The key idea behind these algorithms is a decomposition of each image block into a polyharmonic component and a residual. The polyharmonic component in this talk is an approximate solution to Poisson's equation with Neumann boundary conditions, which means that it is a smooth predictor of the original image block only using the image gradient information across the block boundary. Thus the residual – obtained by removing the polyharmonic component from the original image block – has approximately zero gradient across the block boundary which gives rise to fast-decaying DCT coefficients, which in turn lead to more efficient compression-decompression algorithms for the same bitrates. Our numerical experiments objectively and subjectively demonstrate the superiority of PHLCT over the JPEG Baseline method and the improvement of JPEG-compressed images when decompressed by PPHLCT.

If time permits, I will also discuss my idea of using the Helmholtz equation instead of the Poisson/Laplace equations for images containing oscillatory textures.

3.21 Reconstruction of Sparse Exponential Polynomials from Samples

Tomas Sauer (Universität Passau, DE)

Prony's method, in its various concrete algorithmic realisations, is concerned with the reconstruction of a sparse exponential sum from integer samples. In several variables, the reconstruction is based on finding the variety for a zero dimensional radical ideal. If one replaces the coefficients in the representation by polynomials, i.e., tries to recover sparse exponential polynomials, the zeros associated to the ideal have multiplicities attached to them. Indeed, there is a precise and fairly explicit relationship between the coefficients in the exponential polynomial and the multiplicity spaces of zeros.

3.22 Anisotropic Surface Interpolation

Carola-Bibiane Schönlieb (University of Cambridge, GB)

Joint work of M. Benning, C. Brune, J. Lellmann, S. Masnou, J.-M. Morel, S. Parisotto, C.-B. Schönlieb

In this talk we discuss the use of anisotropic total variation regularisation for interpolating highly structured functions. We will motivate the modelling, discuss its numerical solution and show applications to digital elevation maps and limited angle tomography.

3.23 Remote Sensing Data Compression

Joan Serra Sagristà (Autonomus University of Barcelona, ES)

Remote sensing data compression is an active area of research and has seen some critical developments recently. We will first motivate the need for efficient data transmission – where data compression plays a significant role – in current and future Earth Observation missions. Then we will introduce some of our own recent developments in this field.

3.24 Inpainting-Based Image Compression: State-of-the-Art, Challenges, and Connections

Joachim Weickert (Universität des Saarlandes, DE)

Inpainting-based approaches for lossy image compression store only a small, carefully selected subset of the image data. In the decoding phase, the missing data are reconstructed by a suitable inpainting process. This can be achieved e.g. by interpolation with partial differential equations (PDEs) or by copying information from patches in other image regions. Such codecs allow a very intuitive interpretation, and first experiments show their advantages for high compression rates where they can beat even advanced transform-based methods. However, much more research is needed to turn them into viable codecs. In this talk we will discuss the state-of-the-art, identify challenging problems, and mention relations to other fields. We will see that these problems can only be solved, if different communities with expertise on inpainting, approximation theory, optimisation, numerical analysis, and coding collaborate in an interdisciplinary way.

3.25 An Approximated PDHG Method for Total Variation Wavelet Inpainting

Hao-Min Zhou (Georgia Institute of Technology – Atlanta, US)

Joint work of Xiaojing Ye, Hao-Min Zhou

The primal-dual hybrid gradient (PDHG) algorithm has been successfully applied to a number of total variation (TV) based image reconstruction problems for fast numerical solutions. We show that PDHG can also effectively solve the computational problem of image inpainting in wavelet domain, where high quality images are to be recovered from incomplete wavelet coefficients due to lossy data transmission. In particular, we propose an approximated PDHG algorithm to tackle the non-orthogonality of Daubechies 7-9 wavelet which is widely used in practice. We show that this approximated version essentially alters the gradient descent direction in the original PDHG algorithm, but eliminates its orthogonality restriction and retains low computational complexity.

4 Panel discussions

4.1 Panel Discussion on Inpainting-Based Image Compression

Thomas Pock (TU Graz, AT)

The panel discussion took place on Thursday, 17.11.2016, 20:00–21:00. Its goal was to discuss open research questions on the interface between image inpainting and image compression. The panel consisted of five panel members: Joachim Weickert, Jalal Fadili, Carola-Bibiane Schönlieb, Christine Guillemot, and Gerlind Plonka-Hoch. The moderator of the panel discussion was Thomas Pock.

The panel discussion started with a question to Joachim Weickert, why it should be a good idea to use image inpainting techniques for image and video compression algorithms. Joachim Weickert referred to the properties of the human visual system, which is mainly guided by edges and textures. Compressing this information in a suitable way, it might therefore be a good idea to reconstruct the missing image information be means of image inpainting. Furthermore, people working in image inpainting need to understand and solve the same kind of problems as people working in image compression. Gerlind Plonka-Hoch pointed out that this was actually always the goal of the wavelet transform or image adaptive variants that have been proposed in the past.

When asked if image inpainting based techniques will lead to new codecs, Christine Guillemot pointed out that it is very difficult to establish new standards for image and video codecs. First, people who have worked for more than 30 years on codecs will not easily accept a new technique, second, existing codecs might be "good enough", and third, companies who have already created dedicated hardware for established codecs will try to promote their existing algorithms simply to make more money. This is in particular true for the JPEG2000 standard which on the mass market was never able to become a substitute for its predecessor JPEG. Christine also expressed her feeling that inpainting-based compression algorithms will probably be relevant only for certain niche applications.

At this point, Jalal Fadili told the story of a very specialised imaging device that is used in the Herschel space probe. He was part of the team developing a compressed sensing technique for this device. Although the compressed sensing technique would have led to improved image quality, JPEG2000 was used in the end, mainly due to software reasons.

To account for that, Joachim Weickert remarked that inpainting-based image compression should be as simple as possible (e.g. linear, homogeneous diffusion) in order to allow for fast decoding. Joachim Weickert also pointed out that each neuron in the human brain is doing only very simple operations. To understand how the complexity of natural images can be realised by very simple and local operations will be a major future research questions.

Another topic that was discussed in the panel was video compression. During the week we saw very interesting and impressive results for video inpainting but almost no work has been done to use such techniques for video compression. One of the reasons is the high computational complexity of video inpainting algorithms that hinders those algorithms from real-time applications. Moreover, the human perception of videos is often quite counter-intuitive and hence it is difficult to define suitable quality measurements which take these aspects into account.

Finally, the panel discussed the role of quality measurements for image compression. It turns out that it is still completely unclear which performance measure is the most relevant.

In fact, image compression might be even task-dependent. Carola Schönlieb pointed out that when trying to find a lost object in an image or identifying a certain person in a video, one might be interested in certain details which could have been lost when compressing the image or video. Moreover, in medical applications, CT or MRI images are actually measuring certain physical quantities and medical doctors will in fact base their diagnoses on the image content. Hence one has to be even more careful when applying image compression to such images.

In general it turned out that the panel discussion was a very good idea and it led to vital discussions between the panel members and the audience. The discussion showed that there are much more open questions than solutions and that bringing together different communities in this seminar was actually a very good idea.

Participants

- Sarah Andris
Universität des Saarlandes, DE
- Johannes Ballé
New York University, US
- Zakaria Belhachmi
University of Mulhouse, FR
- Aurélien Bourquard
MIT – Cambridge, US
- Eva-Maria Brinkmann
Universität Münster, DE
- Dorin Bucur
University Savoie Mont Blanc – Le Bourget-du-Lac, FR
- Martin Buhmann
Universität Gießen, DE
- Gene Cheung
National Institute of Informatics – Tokyo, JP
- Nira Dyn
Tel Aviv University, IL
- Gabriele Facciolo
ENS – Cachan, FR
- Jalal Fadili
Caen University, FR
- Irena Galic
University of Osijek, HR
- Yann Gousseau
Telecom ParisTech, FR
- Christine Guillemot
INRIA – Rennes, FR
- Laurent Hoeltgen
BTU Cottbus, DE
- Armin Iske
Universität Hamburg, DE
- Claire Mantel
Technical University of Denmark – Lyngby, DK
- Simon Masnou
University Claude Bernard – Lyon, FR
- Peter Ochs
Universität Freiburg, DE
- Pascal Peter
Universität des Saarlandes, DE
- Gerlind Plonka-Hoch
Universität Göttingen, DE
- Thomas Pock
TU Graz, AT
- Daniela Rosca
TU of Cluj-Napoca, RO
- Naoki Saito
Univ. of California – Davis, US
- Tomas Sauer
Universität Passau, DE
- Carola-Bibiane Schönlieb
University of Cambridge, GB
- Joan Serra Sagristà
Autonomus University of Barcelona, ES
- Joachim Weickert
Universität des Saarlandes, DE
- Hao-Min Zhou
Georgia Institute of Technology – Atlanta, U