Report from Dagstuhl Seminar 11471

Efficient Algorithms for Global Optimisation Methods in Computer Vision

Edited by

Andrés Bruhn[1], Thomas Pock[2], and Xue-Cheng Tai[3]

1 **Universität des Saarlandes, DE,** bruhn@mmci.uni-saarland.de
2 **TU Graz, AT,** pock@icg.tugraz.at
3 **University of Bergen, NO and Nanyang Technology University, SG,** tai@mi.uib.no

Abstract

This report documents the program and the results of Dagstuhl Seminar 11471 *Efficient Algorithms for Global Optimisation Methods in Computer Vision*, taking place November 20–25 in 2011. The focus of the seminar was to discuss the design of efficient computer vision algorithms based on global optimisation methods in the context of the *entire* design pipeline. Since there is no such conference that deals with all aspects of the design process – *modelling*, *mathematical analysis*, *numerical solvers*, and *parallelisation* – the seminar aimed at bringing together researchers from computer science and mathematics covering all four fields.

Seminar 20.– 25. November, 2011 – www.dagstuhl.de/11471

1998 ACM Subject Classification I.4 Image Processing and Computer Vision, G.1.6 Global Optimization, G.2.1 Combinatorial Algorithms, F.2.1 Analysis of Algorithms and Problem Complexity, G.1.0 Parallel Algorithms.

Keywords and phrases Computer Vision, Modelling, Mathematical Foundations, Data Structures, Efficient Algorithms, Parallel Computing

Digital Object Identifier 10.4230/DagRep.1.11.66

1 Executive Summary

Andrés Bruhn
Thomas Pock
Xue-Cheng Tai

Most of the leading algorithms in computer vision are based on global optimisation methods. Such methods compute the solution of a given problem as minimiser of suitable cost functional that penalises deviations from previously made assumptions and integrates them over the entire image domain. While their transparent modelling allows for excellent results in terms of quality for many fundamental computer vision tasks such as *motion estimation*, *stereo reconstruction*, *image restoration*, *shape matching*, and *object segmentation*, the minimisation of these cost functional often leads to optimisation problems that are mathematically challenging and computationally expensive.

In the last decade, this fact has triggered a variety of different research directions that try to satisfy the needs of an ever increasing resolution in image and video data as well as the strong real-time demands of industrial applications. These research directions can be roughly divided into four fields that correspond to the different stages of the algorithmic design pipeline:

- *Modelling* (suitable priors, continuous vs. discrete, hybrid approaches)
- *Mathematical Analysis* (convex vs. non-convex, error bounds, well-posedness)
- *Numerical Solvers* (recent techniques, trends in numerical/combinatorial optimisation)
- *Parallelisation* (GPUs, mutli-core, cluster systems, FPGAs)

Since there are no conferences that address all four fields, the goal of this seminar was to identify and address open questions in these four fields in the context of the *entire* design pipeline. To this end, it brought together computer scientists and mathematicians from all stages. Apart from stimulating interdisciplinary research and establishing close collaborations between the different fields by scheduling plenty of time for discussions, the ultimate goal of the seminar was to develop more precise and more efficient algorithms that are conceptually well designed, mathematically well understood and from which all parts are chosen carefully such that they harmonise with each other.

Further aims of the seminar were to establish suitable benchmarks for measuring the performance of each of the stages (model precision, optimisation accuracy, numerical efficiency, parallelisability) and to derive general conceptual guidelines for the design of efficient algorithms that are applicable to a broad class of key problems in computer vision.

The seminar was conducted in a conference style, where every participant gave a talk of about 20 to 25 minutes. There was much time for extensive discussions – directly after the talks, in dedicated working groups and in the evenings. As documented by the very positive evaluation and the detailed summaries of the working groups, there was a very open and constructive atmosphere. In particular the people appreciated the integration of young researchers in the seminar schedule and the gain of new cross-disciplinary insights from talks and discussions of other participants. While the first observation reflects the fact that the field of efficient algorithms in computer vision is relatively new, the second aspect demonstrates the success of the seminar to bring together people from different communities and establish new ties between the fields. At the moment that this report is written it is very difficult to identify new fundamental issues of efficient algorithms in computer vision. However, lots of interesting aspects were discussed during the talks and in the discussion groups, and many participants established collaborations with people from other fields. Thus we believe this seminar served as an excellent basis to inspire and trigger novel developments in the design of efficient algorithms for global optimisation problems in computer vision.

Finally, it should be mentioned that there was a huge consensus among the participants that there should be a follow-up event to this seminar in the upcoming years. For the current seminar, there will be edited post-proceedings in the Springer LNCS series which gives all participants the opportunity to summarise results from the seminar, discuss open questions and present recent research in the field. The deadline for submission is end of March 2012.

2 Table of Contents

3 Overview of Talks

3.1 Simultaneous Convex Optimization of Regions and Region Parameters in Image Segmentation Models

Egil Bae (University of Bergen, NO)

Joint work of Egil Bae, Jing Yuan, Xue-Cheng Tai

This work develops a convex optimization framework for image segmentation models, where both the unknown regions and parameters describing each region are part of the optimization process. Convex relaxations and optimization algorithms are proposed, which produce results that are independent from the initializations and closely approximate global minima. We focus especially on problems where the data fitting term depends on the mean or median image intensity within each region. We also develop a convex relaxation for the piecewise constant Mumford-Shah model, where additionally the number of regions is unknown. The approach is based on optimizing a convex energy potential over functions defined over a space of one higher dimension than the image domain.

3.2 Multi-label Segmentation with Constraints on Geometry and Complexity

Yuri Boykov (Middlesex College – Ontario, CA)

Joint work of Andrew Delong, Lena Gorelick, Anton Osokin, Hossam Isack, Yuri Boykov

Existing segmentation methods rutinely apply mixture models to objects containing spatially distinct regions, each with a unique color/texture model. Since mixture models assume i.i.d. distribution of intensities, the spatial distribution of colors within complex objects is ignored, which significantly limits the robustness and the accuracy of segmentation results. Despite, mixture models are widely used in segmentation because they easily integrate into standard optimization techniques as simple unary potentials.

We discuss segmentation models corresponding to more general non-i.i.d. appearances and propose the corresponding optimization methods. Sub- labels are used for object parts with additional constraints needed to describe the distribution of parts. We focus on two types of generic constraints. The first type models geometric relations between the parts (e.g. “inclusion”, “exclusion”, etc. – useful in bio-medical applications due to anatomy). The second type models complexity of appearance (i.e. MDL principle – useful for both photo and medical imagery). These segmentation models allow either global optimization or sufficiently good approximations.

The talk is based on materials from four recent papers

- http://www.csd.uwo.ca/faculty/yuri/Abstracts/emmcvpr11-sabs.html
- http://www.csd.uwo.ca/faculty/yuri/Abstracts/ijcv10_lc-abs.html
- http://www.csd.uwo.ca/faculty/yuri/Abstracts/iccv11-abs.html
- http://www.csd.uwo.ca/faculty/yuri/Abstracts/iccv09-abs.html

3.3 Total Generalized Variation: A Convex Model for Piecewise Smooth Images

Kristian Bredies (Universität Graz, AT)

Joint work of Kristian Bredies, Martin Holler
Main reference K. Bredies, M. Holler, "Artifact-free JPEG decompression with total generalized variation," Int'l Joint Conf. on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISAPP), 2012.
A preprint is available as SFB-Report No. 2011-029 from the SpezialForschungsBereich F32.
URL http://math.uni-graz.at/mobis/publications/SFB-Report-2011-029.pdf

We introduce and study the total generalized variation functional (TGV) which constitutes a well-suited convex model for piecewise smooth images. It comprises exactly the functions of bounded variation but is, unlike total-variation based functionals, also aware of higher-order smoothness. The feasibility and the effect of regularization with TGV is demonstrated. In particular, we address the design of global optimization methods with TGV penalty. Furthermore, the framework is applied to the reconstruction of color images from inexact and incomplete data: We propose a TGV-based approach for the reduction of typical decompression artifacts in highly compressed JPEG images and show numerical results.

3.4 Completely Convex Formulation of the Chan-Vese Image Segmentation Model

Xavier Bresson (City University – Hong Kong, HK)

The active contours without edges model of Chan and Vese is a popular method for computing the segmentation of an image into two phases. The minimization problem is non-convex even when the optimal region constants are known a priori. In [Chan-Esedoglu-Nikolova 2006], authors provided a method to compute global minimizers by showing that solutions could be obtained from a convex relaxation.

In this talk, we propose a convex relaxation approach to solve the case in which both the segmentation and the optimal constants are unknown for two phases and multiple phases. In other words, we propose a convex relaxation of the popular K-means algorithm. Although the proposed relaxation technique is not guaranteed to find exact global minimizers of the original problem, our experiments show that our method computes tight approximations of the optimal solutions.

3.5 Efficient Video Analysis Tools towards Large Scale Unsupervised Learning

Thomas Brox (Universität Freiburg, DE)

Joint work of Thomas Brox, Narayanan Sundaram, Peter Ochs, Jitendra Malik
Main reference T. Brox, J. Malik, "Object segmentation by long term analysis of point trajectories," European Conference on Computer Vision (ECCV), 2010.
URL http://dx.doi.org/10.1007/978-3-642-15555-0_21

Learning tasks such as object detection require access to more training data than what is available today. One way to reduce the burden of manual image annotation is crowd sourcing; an alternative one is learning from videos. In particular, I will present video analysis tools that allow for long term analysis of videos by means of point trajectories. A clustering of these point trajectories enables automatic segmentation of moving objects, thereby providing annotation without any user interaction. Application of these tools to large amounts of data requires efficient algorithms that can make use of cheap parallel hardware, such as GPUs.

3.6 Continuous Limits of Discrete Perimeters

Antonin Chambolle (Ecole Polytechnique – Palaiseau, FR)

Main reference A. Chambolle, A. Giacomini, L. Lussardi, "Continuous limits of discrete perimeters," ESAIM: M2AN, 44(2), pp. 207–230, 2010.
URL http://dx.doi.org/10.1051/m2an/2009044

We have considered in this talk the continuous limit of various discrete approximations of the perimeter and/or total variation. After a discussion on how good can an isotropic approximation be, we have studied a class of discrete convex functionals which satisfy a (generalized) coarea formula. These correspond to approximations with discrete submodular functions, of not necessarily isotropic but translational invariant—homogeneous—perimeters. In these cases, indeed, perimeter minimization problems can be solved exactly with polynomial algorithms. We have described the variational limit, as the size of the pixels go to zero, of such perimeters, and shown in particular that they always correspond to a crystalline interfacial energy (which has an explicit form). A question which is still unclear is whether the limits of arbitrary submodular interactions form a larger set than the limits of approximations with pairwise interactions (which can be solved by max flow/min cut algorithms on a graph).

3.7 A Variational Approach for Exact Histogram Specification

Raymond Chan (Chinese University of Hong Kong, HK)

We focus on exact histogram specification when the input image is quantified. The goal is to transform this input image into an output image whose histogram is exactly the same as a prescribed one. In order to match the prescribed histogram, pixels with the same intensity level in the input image will have to be assigned to different intensity levels in the output image. An approach to classify pixels with the same intensity value is to construct a strict ordering on all pixel values by using auxiliary attributes. Local average intensities and wavelet coefficients have been used by the past as the second attribute. However, these methods cannot enable strict-ordering without degrading the image.

In this paper, we propose a variational approach to establish an image preserving strict-ordering of the pixel values. We show that strict-ordering is achieved with probability one. Our method is image preserving in the sense that it reduces the quantization noise in the input quantified image. Numerical results show that our method gives better quality images than the existing methods.

3.8 Variational Methods for Image Reconstruction in Magnetic Resonance Imaging

Christian Clason (Universität Graz, AT)

Joint work of Christian Clason, Florian Knoll, Kristian Bredies
Main reference F. Knoll, C. Clason, K. Bredies, M. Uecker, R. Stollberger, "Parallel imaging with nonlinear reconstruction using variational penalties," Magnetic Resonance in Medicine, 67(1), pp. 34–41, 2012.
URL http://dx.doi.org/10.1002/mrm.22964

Acquisition speed is currently the main limiting factor for magnetic resonance imaging. One well established technique for scan time reduction acquires only a subset of the necessary data, but with multiple independent receivers, each of which has a different spatial profile which needs to be estimated. Standard reconstruction methods in current clinical practice are linear algebraic in nature and suffer from residual artifacts due to missing information and incorrectly estimated modulations. This talk demonstrates that a variational approach including penalties of total (generalized) variation type can significantly improve the reconstruction quality for accelerated image acquisition.

3.9 Geometrically Consistent Elastic Matching of 3D Shapes: A Linear Programming Solution

Daniel Cremers (TU München, DE)

Joint work of Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers
Main reference T. Windheuser, U. Schlickewei, Frank R. Schmidt, D. Cremers, "Geometrically consistent elastic matching of 3D shapes: a linear programming solution," IEEE Int'l Conf. on Computer Vision (ICCV), 2011.
URL http://cvpr.in.tum.de/_media/spezial/bib/wssc_iccv11.pdf

We propose a method for computing a geometrically consistent and spatially dense matching between two 3D shapes. Rather than mapping points to points we consistently match infinitesimal surface patches in a manner which minimizes an elastic thin-shell energy, thereby minimizing the physical energy required to deform one shape into the other. We observe that any conceivable matching corresponds to a codimension-2 surface in the 4D product space spanned by the two shapes. Subsequently we show that a discrete formulation of this problem is given by a binary linear program whose relaxed version can be solved efficiently in a globally optimal manner.

Experimental results demonstrate both qualitatively and quantitatively that the proposed LP relaxation allows to compute consistent and accurate matchings which reliably put into correspondence highly deformed and articulated 3D shapes consisting of thousands of triangles.

3.10 Convex Relaxations for Vectorial Multilabel Problems and the Total Curvature

Bastian Goldluecke (TU München, DE)

We discuss two (unrelated) convex relaxation techniques we recently introduced.

The first is for multilabel problems with a large number of vectorial labels, where we can drastically reduce the dimensionality of the problem by exploiting the special structure of the space. The second relaxation is for a novel continuous regularizer we call the Total Curvature of a function, which allows to use this notion of curvature as a replacement for total variation in image processing applications.

3.11 Are Extrapolation Methods Useful for Diffusion-based Image Analysis?

Sven Grewenig (Universität des Saarlandes, DE)

Diffusion methods are widely used in image processing and computer vision, e.g. for image denoising, interpolation, optic flow, and stereo reconstruction. There are two popular ways to implement them: Explicit finite difference schemes are simple but become inefficient due to severe time step size restrictions, while semi-implicit schemes are more efficient but require to solve large linear systems of equations. An efficient alternative is the recently proposed Fast Explicit Diffusion (FED) scheme that allows very large explicit time steps.

However, all three methods have first order accuracy in time. In this talk, we are going to consider more accurate time extrapolation methods based on semi-implicit and FED schemes. To this end, we present two stable FED extrapolation methods and analyse them with respect to their accuracy and efficiency for parabolic image enhancement as well as elliptic image interpolation problems.

3.12 Nonconvex TV^q-Models in Image Restoration: Analysis and a Trust-Region Regularization Based Superlinearly Convergent Solver

Michael Hintermueller (HU Berlin, DE)

Joint work of Michael Hintermueller, Tao Wu
Main reference M. Hintermüller, T. Wu, "Nonconvex TV^q-Models in Image Restoration: Analysis and a Trust-Region Regularization Based Superlinearly Convergent Solver," SFB-Report No. 2011-026, SpezialForschungsBereich F32, 2011.
URL http://math.uni-graz.at/mobis/publications/SFB-Report-2011-026.pdf

A nonconvex variational model is introduced which contains the l^q-norm, with $0 <= q <= 1$, of the gradient of the image to be reconstructed as the regularization term together with a least-squares type data fidelity term which may depend on a possibly spatially dependent weighting parameter. Hence, the regularization term in this functional is a nonconvex compromise between the minimization of the support of the reconstruction and the classical convex total variation model. In the discrete setting, existence of a minimizer is proven, a Newton-type solution algorithm is introduced and its global as well as locally superlinear convergence are established. The potential indefiniteness (or negative definiteness) of the Hessian of the objective during the iteration is handled by a trust region based regularization scheme. The performance of the new algorithm is studied by means of a series of numerical tests. For the associated infinite dimensional model an existence result based on the weakly lower semicontinuous envelope is given and its relation to the original problem is discussed.

3.13 Unsupervised Multiphase Segmentation and Its Regularity

Sung Ha Kang (Georgia Institute of Technology, US)

Variational models have been widely studied for imaging applications in recent years, and one of the most well known model is the Mumford-Shah functional was introduced in the late 1980s. In this talk, we discuss mainly two models. The first model is infinite perimeter segmentation model which is a variant of piecewise constant Mumford-Shah model (typically referred to as Chan-Vese model). It allows fine details and irregularities of the boundaries while denoising additive Gaussian noise. The second model is a multiphase segmentation model which automatically chooses a favorable number of phases as it segments the image. The model is discussed, connection is made with a regularized k-means (for 1-dimentional data clustering application), and new regularity analysis results are presented.

3.14 Metric Geometry In Action

Ron Kimmel (Technion – Haifa, IL)

I will talk about methods imported from metric geometry that help in comparing and matching shapes. Picking up the right metric for the task at hand plays an important role in shape analysis. I will review some tools for that goal.

3.15 Efficient Imaging Algorithms on Many-Core Platforms

Harald Koestler (Universität Erlangen-Nürnberg, DE)

Main reference H. Koestler, "A Multigrid Framework for Variational Approaches in Medical Image Processing and Computer Vision," Dr. Hut Verlag, Muenchen, 2008.

Today it is possible to achieve real-time performance for medium to large size imaging applications even on usual desktop computers or small clusters. For numerical algorithms like multigrid that work on structured grids GPUs show good speedup factors compared to current CPUs mainly due to their higher memory bandwidth. But also other imaging applications like sparse coding profit from GPUs.

Besides performance results for high dynamic range compression and image denoising we will also present some newer ideas for efficient image deblurring algorithms.

3.16 Maxflow-based Inference for Functions with High-order Terms

Vladimir Kolmogorov (IST Austria – Klosterneuburg, AT)

The maxflow technique is an important tool for minimizing functions with pairwise terms. It can compute a global minimum if all terms are submodular, and for general pairwise terms it can identify a part of an optimal solution by solving the roof duality relaxation. I will consider extensions of these techniques to higher-order terms. Previously proposed extensions converted the function to a sum of pairwise terms by introducing auxiliary variables. I will discuss a more direct approach that does not require such a conversion.

3.17 Volumetric Nonlinear Anisotropic Diffusion on GPUs

Arjan Kuijper (FhG IGD – Darmstadt, DE)

Joint work of Andreas Schwarzkopf, Thomas Kalbe, Chandrajit Bajaj, Arjan Kuijper, Michael Goesele
Main reference T. Kalbe, A. Schwarzkopf, M. Goesele, A. Kuijper, C. Bajaj, "Volumetric Nonlinear Anisotropic Diffusion on GPUs," Scale Space and Variational Methods in Computer Vision (SSVM), 2012.
URL http://dx.doi.org/10.1007/978-3-642-24785-9_6

We present an efficient implementation of volumetric nonlinear anisotropic image diffusion on modern programmable graphics processing units (GPUs). We avoid the computational bottleneck of a time consuming eigenvalue decomposition in R3.

Instead, we use a projection of the Hessian matrix along the surface normal onto the tangent plane of the local isodensity surface and solve for the remaining two tangent space eigenvectors. We derive closed formulas to achieve this resulting in efficient GPU code. We show that our most complex volumetric nonlinear anisotropic diffusion gains a speed up of more than 600 compared to a CPU solution.

3.18 Enforcing Topological Constraints in Random Field Image Segmentation

Christoph Lampert (IST Austria – Klosterneuburg, AT)

Joint work of Chao Chen, Daniel Freedman, Christoph H. Lampert
Main reference C. Chen, D. Freedman, C. H Lampert, "Enforcing topological constraints in random field image segmentation," IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2011.
URL http://dx.doi.org/10.1109/CVPR.2011.5995503

We introduce TopoCut, a new way to integrate knowledge about topological properties (TPs) into random field image segmentation model. Instead of including TPs as additional constraints during minimization of the energy function, we devise an efficient algorithm for modifying the unary potentials such that the resulting segmentation is guaranteed with the desired properties.

Our method is more flexible in the sense that it handles more topology constraints than previous methods, which were only able to enforce pairwise or global connectivity. In particular, our method is very fast, making it for the first time possible to enforce global topological properties in practical image segmentation tasks.

3.19 Image Restoration under Hessian Matrix-Norm Regularization

Stamatis Lefkimmiatis (EPFL – Lausanne, CH)

Joint work of Stamatis Lefkimmiatis; Aurelien Bourquard; Michael Unser
Main reference S. Lefkimmiatis, A. Bourquard, M. Unser, "Hessian-based norm regularization for image restoration with biomedical applications," IEEE Trans. on Image Processing, 2012, *to appear*.
URL http://dx.doi.org/10.1109/TIP.2011.2168232

We will discuss about non-quadratic Hessian-based regularization methods that can be effectively used for inverse imaging problems, in a variational framework. Motivated by the great success of the total-variation (TV) functional in the image-restoration setting, we will present an extension that also includes second-order differential operators. Specifically, we will focus on second-order functionals that involve matrix norms of the Hessian operator.

Their definition is based on an alternative interpretation of TV that relies on mixed norms of directional derivatives. We will show that the resulting regularizers retain some of the most favorable properties of TV, namely, convexity, homogeneity, rotation, and translation invariance, while dealing effectively with the staircase effect. We will further present an efficient minimization scheme for the corresponding objective functions. The employed algorithm is of the iterative-reweighted-least-squares (IRLS) type and results from a majorization-minimization approach. It also relies on a problem-specific preconditioned conjugate-gradient method (PCG) that makes the overall minimization scheme very attractive, since it can be effectively applied to large images in a reasonable computational time.

3.20 Optimality Bounds and Optimization for Image Partitioning

Jan Lellmann (Universität Heidelberg, DE)

Joint work of Jan Lellmann, Frank Lenzen, Christoph Schnörr
Main reference J. Lellmann, F. Lenzen, C. Schnörr, "Optimality bounds for a variational relaxation of the image partitioning problem," arXiv:1112.0974v1 [cs.CV]
URL http://arxiv.org/abs/1112.0974v1

Variational convex relaxations can be used to compute approximate minimizers of optimal partitioning and multiclass labeling problems on continuous domains. We present recent developments on proving a priori upper bounds for the objective of the relaxed problem, and on the numerical solution of such problems.

3.21 A Variational Approach for the Fusion of Exposure Bracketed Images

Stacey Levine (Duquesne University, US)

Joint work of Stacey Levine, Marcelo Bertalmio

We propose a variational approach for fusing a set of images taken with different exposure times so that optimal information is obtained from each one. The solution is a single image whose details and edges are extracted from a short exposure time image (typically low contrast and noisy) and color information is extracted from a long exposure time image (often suffering from blur due to motion by the camera and/or the subject(s)). The approach is well posed and results compare favorably with the state of the art.

3.22 Total Variation Denoising Using Posterior Expectation

Cecile Louchet (Université d'Orléans, FR)

Total variation image denoising was originally described as a variational method, but it can be interpreted in a Bayesian framework as a Maximum A Posteriori estimate. This maximization aspect is partly responsible for the so-called staircasing effect, i.e. the outbreak of quasi-constant regions separated by sharp edges in the intensity map.

In this talk we propose to transpose this denoising method into an estimation based on the posterior expectation, in order to better account for the global properties of the posterior distribution. Theoretical and numerical results are presented, which demonstrate in particular that images denoised with the proposed scheme do not suffer from the staircasing effect. We then focus on the practical computation of TV-LSE, and propose an MCMC (Monte-Carlo Markov Chain) algorithm whose convergence is carefully analyzed. Both the model and the algorithm are flexible enough to be directly applied to other low-level image processing tasks, such as image deblurring and interpolation.

3.23 Hyperelasticity in Correspondence Problems

Jan Modersitzki (Universität zu Lübeck, DE)

Image registration is one of the challenging problems in image processing. Given two images taken for example at different times, from different devices or perspectives, the goal is to determine a reasonable transformation, such that a transformed version of one of the images is similar to the second one.

The problem is typically phrased in a variational setting for the wanted transformation field where meaningful regularization is a key issue. In this talk, we present an extension of the well-known linear elastic regularizer. Hyperelasticity is used and enables the modelling of large and non-linear transformations and ensures diffeomorphic mappings. The price to be paid is a non-convex but polyconvex objective function. We also Moreover, a stable and efficient numerical scheme is challenging.

In this talk, we present an introduction to hyperelastic image registration with applications from cardiac PET imaging. We introduce a new numerical framework.

This framework is based on the discretize then optimize paradigm and uses a sophisticated computation of the discrete analogues of the three invariants of the transformation tensor. We show several numerical examples that illustrate the potential of the hyperelastic regularizer as well as some application where hyperelastic regularization is mandatory.

3.24 Fast Regularization of Matrix-Valued Images

Guy Rosman (Technion – Haifa, IL)

Joint work of Guy Rosman, Yu Wang, Xue-Cheng Tai, Ron Kimmel, Alfred M. Bruckstein
Main reference G. Rosman, Y. Wang, X.-C. Tai, R. Kimmel, A. M. Bruckstein, "Fast regularization of matrix-valued images," UCLA CAM Report 11-87, December 2011.
URL ftp://ftp.math.ucla.edu/pub/camreport/cam11-87.pdf

Regularization of matrix-valued data is important in many fields, such as medical imaging, motion analysis and scene understanding, where accurate estimation of diffusion tensors or rigid motions is crucial for higher-level computer vision tasks. In this report we describe a novel method for efficient regularization of matrix group-valued images.

Using the augmented Lagrangian framework we separate the total-variation regularization of matrix-valued images into a regularization and projection steps, both of which are fast and parallelizable. Furthermore we extend our method to a high-order regularization scheme for matrix-valued functions.

We demonstrate the effectiveness of our method for denoising of several group-valued image types, with data in SO(n), SE(n), and SP D(n), and discuss its convergence properties.

3.25 Massively Parallel Multigrid for Image Processing

Ulrich Ruede (Universität Erlangen-Nürnberg, DE)

Multigrid is among the most efficient solvers for global problems, since the hierarchical system of coarser grids provide an elegant means to exchange information in the problem domain. Many variational image processing problems lead to elliptic differential equations that need this kind of data exchange. While the coarser grids are therefore essential for fast global solvers, it may become a problem for the parallelization. In the presentation we will demonstrate that a careful implementation can nevertheless show excellent scalability and provide powerful solvers also for very large image data sets.

3.26 Quadrature Rules, Discrepancies and Variational Dithering

Gabriele Steidl (TU Kaiserslautern, DE)

This paper deals with continuous-domain quantization which aims to create the illusion of a gray-value image by appropriately distributing black dots. For lack of notation we refer to the process as halftoning which is usually associated with the quantization on a discrete grid. Recently a framework for this task was proposed by minimizing an attraction-repulsion functional consisting of the difference of two continuous, convex functions. One of them describes attracting forces caused by the image gray values, the other one enforces repulsion between dots. In this paper, we generalize this approach by considering quadrature error functionals on reproducing kernel Hilbert spaces (RKHSs) with respect to the quadrature nodes, where we ask for optimal distributions of these nodes. For special reproducing kernels these quadrature error functionals coincide with discrepancy functionals. It turns out that the attraction-repulsion functional appears for a special RKHS of functions on $\mathbb{R}^2$. Moreover, our more general framework enables us to consider optimal point distributions not only in $\mathbb{R}^2$ but also on the torus $\mathbb{T}^2$ and the sphere $\mathbb{S}^2$. For a large number of points the computation of such point distributions is a serious challenge and requires fast algorithms. To this end, we work in RKHSs of bandlimited functions on $\mathbb{T}^2$ and $\mathbb{S}^2$. Then, the quadrature error functional can be rewritten as a least squares functional. We propose a nonlinear conjugate gradient method to compute a minimizer of this functional and show that each iteration step can be computed in an efficient way by fast Fourier transforms at nonequispaced nodes on the torus and the sphere. Numerical examples demonstrate the good quantization results obtained by our method.

3.27 Generalized Ordering Constraints for Multilabel Optimization

Evgeny Strekalovskiy (TU München, DE)

Joint work of Evgeny Strekalovskiy, Daniel Cremers
Main reference E. Strekalovskiy, D. Cremers, "Generalized ordering constraints for multilabel optimization," IEEE Int'l Conf. on Computer Vision (ICCV), pp. 2619–2626, 2011.
URL http://dx.doi.org/10.1109/ICCV.2011.6126551

We propose a novel framework for imposing label ordering constraints in multilabel optimization. In particular, label jumps can be penalized differently depending on the jump direction. In contrast to the recently proposed MRF-based approaches, the proposed method arises from the viewpoint of spatially continuous optimization. It unifies and generalizes previous approaches to label ordering constraints:

Firstly, it provides a common solution to three different problems which are otherwise solved by three separate approaches. We provide an exact characterization of the penalization functions expressible with our approach.

Secondly, we show that it naturally extends to three and higher dimensions of the image domain. Thirdly, it allows novel applications, such as the convex shape prior. Despite this generality, our model is easily adjustable to various label layouts and is also easy to implement. On a number of experiments we show that it works quite well, producing solutions comparable and superior to those obtained with previous approaches.

3.28 Parallel Preconditioners for GPU-Multigrid Solvers

Robert Strzodka (MPI für Informatik – Saarbrücken, DE)

Joint work of Dominik Göddeke, Robert Strzodka
Main reference D. Göddeke, R. Strzodka, "Mixed precision GPU-multigrid solvers with strong smoothers," in Jack J. Dongarra, David A. Bader, and Jakub Kurzak (eds.): Scientific Computing with Multicore and Accelerators, pp. 131–147, CRC Press, 2010.

GPUs perform best on regular, independent work-loads while effective preconditioners ask for complex, serially coupled computations. The extreme cases of best hardware performance on slowly converging numerical schemes or quickly converging schemes with slow serial execution are poor choices. The talk discusses how to find the middle ground by challenging the hardware with more complex data dependencies and at the same time relaxing purely serial dependencies with parallel variants. For structured matrices, we can now solve very ill-conditioned linear equation systems that were intractable with GPU hardware before.

References

1 D. Göddeke and R. Strzodka. Cyclic reduction tridiagonal solvers on GPUs applied to mixed precision multigrid. *IEEE Transactions on Parallel and Distributed Systems (TPDS), Special Issue: High Performance Computing with Accelerators*, 22(1):22-32, January 2011

3.29 Image Segmentation by Iteratively Optimization of Multiphase Multiple Piecewise Constant Model and Four-Color Relabeling

Wenbing Tao (Huazhong University of Science & Technology, CN)

An iteratively unsupervised image segmentation algorithm is developed, which is based on our proposed multiphase multiple piecewise constant (MMPC) model and its graph cuts optimization. The MMPC model use multiple constants to models each phase instead of one single constant used in Chan and Vese (CV) model and cartoon limit so that heterogeneous image object segmentation can be effectively deal with.

We show that the multiphase optimization problem based on our proposed model can be approximately solved by graph cuts methods. The Four-Color theorem is used to relabel the regions of image after every iteration, which makes it possible to represent and segment an arbitrary number of regions in image with only four phases. Therefore, the computational cost and memory usage are greatly reduced. The comparison with some typical unsupervised image segmentation methods using a large number of images from the Berkeley Segmentation Dataset demonstrates the proposed algorithm can effectively segment natural images with good performance and acceptable computational time.

3.30 Nonlocal Filters for Removing Multiplicative Noise

Tanja Teuber (TU Kaiserslautern, DE)

Joint work of Tanja Teuber, Annika Lang
Main reference T. Teuber, A. Lang, "Nonlocal filters for removing multiplicative noise," Scale Space and Variational Methods in Computer Vision (SSVM), 2012.

We present new nonlocal filters for removing multiplicative noise in images. The considered filters are deduced by weighted maximum likelihood estimation and the occurring weights are defined by a new similarity measure specially designed for comparing data corrupted by multiplicative noise. For the deduction of this measure we analyze a probabilistic measure recently proposed for general noise models by Deledalle et al. The properties of our new measure are examined theoretically as well as by numerical experiments. Moreover, different adaptations are proposed to further improve the results. Throughout the talk, our findings are exemplified for multiplicative Gamma noise. Finally, restoration results are presented to demonstrate the good properties of our new filters.

3.31 Optimization for Pixel Labeling Problems With Structured Layout

Olga Veksler (University of Western Ontario, CA)

Joint work of Olga Veksler, Pedro Felzenszwalb, Xiaoqing Liu
Main reference P. F. Felzenszwalb, O. Veksler, "Tiered scene labeling with dynamic programming," IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2010.
URL http://dx.doi.org/10.1109/CVPR.2010.5540067

Pixel labeling problems are pervasive in computer vision research. In this talk, we discuss optimization approaches for labeling problems which have some structure imposed on the layout of the labels. In other words, the relationships between labels is not arbitrary but has a well defined spatial structure. We will describe two approaches for structured layout scenes.

The first approach is for a more restrictive type of scenes, for which we develop new graph-cut moves which we call order-preserving. The advantage of order preserving moves is that they act on all labels simultaneously, unlike the popular expansion algorithm, and, therefore, escape local minima more easily. The second approach is for a more general type of structured layout scenes and it is based on dynamic programming.

In the second case, the exact minimum can be found efficiently. This is very rare for a 2D labeling problem to have an efficient and global optimizer. For both approaches, our applications include geometric class labeling and segmentation with a shape prior.

3.32 3D Motion Analysis with Stereo Cameras

Andreas Wedel (Daimler AG, DE)

Joint work of Andreas Wedel, Clemens Rabe, Thomas Müller, Uwe Franke
Main reference C. Rabe, T. Müller, A. Wedel, U. Franke, "Dense, robust, and accurate motion field estimation from stereo image sequences in real-time," European Conf. on Computer Vision (ECCV), 2010.
URL http://dx.doi.org/10.1007/978-3-642-15561-1_42

We describe the estimation of the three dimensional motion field of the visible world from stereo image sequences. Our approach is build upon the ideas of "dense scene flow" and "6d vision". It combines dense variational optical flow estimation, including spatial regularization, with Kalman filtering for temporal smoothness and robustness. Using a GPU and an FPGA for the implementation yields a vision-system which is directly applicable in real-world scenarios, like automotive driver assistance systems or in the field of surveillance.

3.33 Discrete Theory and Numerical Analysis of Anisotropic Diffusion Interpolation

Joachim Weickert (Universität des Saarlandes, DE)

Joint work of Joachim Weickert, Sven Grewenig, Markus Mainberger, Martin Welk

Image compression with partial differential equations (PDEs) is based on the idea to store only a few pixels and to perform PDE-based inpainting inbetween. Experiments have shown that edge-enhancing anisotropic diffusion (EED) outperforms other PDEs for this task. However, it creates a number of numerical problems: There is no general energy functional, the process is highly anisotropic, dissipative artifacts may occur, and rotation invariance may be approximated badly. Since the theory of nonnegative matrices cannot be applied, no discrete maximum-minimum principle holds. Moreover, the process can be highly ill-conditioned, and linearisations of the nonlinear system of equations for the inpainting problem lead to unsymmetric system matrices. In my talk I will address these problems and show how some of them can be solved.

3.34 Variational Optical Flow

Manuel Werlberger (TU Graz, AT)

Joint work of Manuel Werlberger; Thomas Pock, Horst Bischof
Main reference M. Werlberger, T. Pock, H. Bischof, "Motion estimation with non-local total variation regularization," IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2010.
URL http://dx.doi.org/10.1109/CVPR.2010.5539945

The presented motion estimation algorithm is based on non-local total variation regularization which allows for integrating low level image segmentation in a unified variational framework. Based on Gestalt principles of grouping areas are combined together and the similarity of motion within these regions is enforced.

For a more robust data fidelity term we propose not to linearize the intensity function as for standard algorithms but to approximate the data term directly with a second order Taylor expansion. This allows for using arbitrary data terms.

3.35 Optimization on Spheres and Stiefel Manifolds, with Applications on Genus-0 Surfaces

Wotao Yin (Rice University – Houston, US)

Joint work of Xianfeng Gu; Lai, Lu Rongjie, Ming Lok; Zaiwen Wen, Wotao Yin

Minimization with orthogonality constraints ($X'X = I$) and/or spherical constraints ($||x||_2 = 1$) has wide applications in polynomial optimization, combinatorial optimization, eigenvalue problems, sparse PCA, p-harmonic flows, 1-bit compressive sensing, matrix rank minimization, etc. These problems are non-convex, and it is numerically expensive to preserve the constraints while improving the objective.

We propose to use a Crank-Nicolson-like update scheme that preserves these constraints while performing descent on the manifold with curvilinear search. The update has a low per-iteration cost compared to those based on projections and geodesics. This talk describes the update formula, analyzes its complexity, and presents the full algorithm with the application of global conformal parameterization of genus-0 surfaces.

The optimization part is joint work with Zaiwen Wen. The conformal parameterization part is joint work with Rongjie Lai, Zaiwen Wen, Xianfeng Gu, and Lok Ming Lu. See http://optman.blogs.rice.edu for slides and preprints.

4 Working Groups

Based on proposals from the seminar participants three main working groups were established that met several times during the time of the seminar:

- Models and Priors
- Convex Methods
- Parallelization

Below we give detailed summaries of topics and outcomes of the different workings groups.

4.1 Models and Priors

The first working group addressed the fundamental question of designing models and priors for imaging problems.

In energy minimization approaches the cost functional often penalizes a certain norm on specific image features such as the image gradient or the residual between the noisy image and the unknown clean image. In the discussion of this working group, the following evolution of different norms was pointed out. Starting with the first work in image processing and computer vision in the eighties, mainly quadratic L^2 norms have been used in the cost functional. The main reason for using quadratic norms is based on the fact, that computing a minimizer of a quadratic cost functional amounts for solving systems of linear equations, which was a tractable problem at this time. With the development of new powerful optimization methods (mainly convex optimization algorithms) in the nineties, L^1 norm-type penalizers became very popular, mainly due to their robustness and edge-preserving properties. Later, it has been observed that so-called pseudo-norms (L^p norms for $1 < p < 1$) perform even better, but due to their non-convexity, these methods require much more sophisticated numerical optimization algorithms. Recently, the so-called L^0 norm became very popular in different branches of image and signal processing. Its basic idea is to find a very sparse representation of a signal in a suitable basis which allows for very accurate reconstructions from very few noisy measurements. In the end of this discussion, the question was raised what should come next, but no clear direction can be identified so far.

Another fundamental question that has been discussed in this working group was the question of higher-order methods, i.e. methods that penalize higher-order derivatives in images. While for almost three decades, methods penalizing first order derivatives are dominating, it has been observed quite early, that higher-order methods provide a much stronger prior for imaging tasks. However, higher-order methods are much more demanding in terms of computational time and very often require the design of specialized numerical optimization algorithms. The consensus of this discussion was, that much more work has to be done into this direction to exploit the known potential of higher order methods.

Finally, the fundamental question of how to design or compute good image priors was discussed. On the one hand, researchers use "deterministic" image priors such as the total variation, which is known to favor piecewise constant images. This works well for images of this kind but clearly does not respect the diversity of natural images. Another approach, which is mainly used in image compression is to use local patch based discrete cosine transforms. This prior works much better for natural images but still does not yield satisfactorily results for many problems. On the other hand researchers try to "learn" good image priors. In recent work an algorithm is proposed, which simply fits a high-dimensional Gaussian mixture model to a large number of patches that have been extracted from an image database. it turns out that indeed a large portion of the structure that is learned coincides with oriented image derivatives or with oscillating patterns that can be found in the basis vectors of the discrete cosine transform. The working group concluded that there are fundamental relationships and that it is very important to investigate these methods.

4.2 Convex Methods

The second discussion group that attracted a large number of participants dealt with the design and mathematical analysis of convex formulations of computer vision problems.

So far, most of the global minimization algorithms for computer vision have been carried out through convexification of the minimization problem. One of the questions discussed is about how far this approach can research and be extended. The majority of the models proposed for image processing and computer vision need to minimize an energy functional that is non-linear and non-convex. Traditionally, algorithms often have problems with local minimizers and dependence on initial values. Some recent research results show that a number of these models can be solved through some convexified problems. The convex problems have global minimizers. These global minimizers are either global minimizers for the original non-conex problems or provide a good approximation. During the discussions, global minimization related to length regularization was discussed. Relationship between several methods recently proposed was reviewed. Potential methods that can overcome some problems related to length approximation were mentioned in the discussion. It is concluded that convexification approach can be used for a large number of imaging and vision problems. What we have discovered so far is just a few of them.

Another crucial issue for global minimization is to extend some of the recently proposed algorithms to higher order regularization techniques. Higher order regularizer are much more complicated. A few techniques have been proposed related to graphs. A number of talks in the workshop have discussed this issue. However, convexification of higher model related to "curvatures" and different norms of the Hessian of the minimization functional is still missing in the continuous setting. Researches and investigations into this direction need to be encouraged.

A third issue for convex global minimization is about efficiency and parallel computing. As imaging and vision problems are getting large with the size of data, it is urgent to improve the efficiency of the global minimization algorithms through different accelerations and parallel implementation techniques. Some good acceleration techniques have been suggested in the literature. However, more research is needed on parallel acceleration techniques.

4.3 Parallelization

The third working group was discussing issues of parallelization, large scale computing and benchmarking.

Due to the strong increase in computational power and intrinsic parallelism, the use of graphics processor units (GPUs) became very popular to accelerate algorithms in the field of computer vision. In this context, papers on GPU optimisation typically report huge speed-ups of up to three orders of magnitude. In this discussion it became evident that in most cases the potential gain should actually be much smaller than reported in the literature – somewhere in the order between a factor two and five. This is due to the fact that most algorithms for global optimisation problems in computer vision are memory-limited such that speed ups are mainly related to an increase in memory bandwidth from CPU to GPU. This in turn means that many CPU algorithms are actually not efficiently implemented, either because researchers are not aware of standard CPU optimisation strategies (SIMD, cache optimisation), or because more time is spent on optimising the GPU code. As a conclusion for this first part, the members of the discussion group agreed that in general absolute run times should be considered instead of relative speed ups which may be highly misleading. Moreover, more research should be done on sequential algorithms before turning towards their parallel implementation.

With respect to large scale computing the discussion showed that the trend seems to go to cluster systems with millions of multi-core CPUs. This has three consequences: On the one hand, huge datasets for benchmarking are required that allow for a reasonable parallelization. Due to the increasing resolution of image acquisition devices such data sets are actually available (e.g. satellite images, video data, photo collections). On the other hand, large scale data sets require a sophisticated memory management based on adaptive or hierarchical strategies. This aspect is new to many researchers in computer vision that worked with small and medium scale problems so far. A third issue is the visualization of large scale data sets which becomes even more important due to the increasing gap between the size of data sets (up to a few petabyte) and the resolution of displays (almost constant at a few megapixels). Seemingly, working on compressed or even adaptively resampled data sets is another challenge that will become more and more important in computer vision. It is an open question how existing algorithms have to be modified such that they allow the processing of such condensed information. As a main conclusion of this second part the group members agreed that the algorithmic design pipeline described in the Dagstuhl proposal will get significantly more complex in the future. In the case of a potential follow up seminar, this should be addressed by inviting also people from the above mentioned fields.

Finally, the aspect of benchmarking the speed and the quality of efficient algorithms was discussed. Apart from the aforementioned problem that parallel implementations are typically compared to poorly optimised sequential schemes, a major issue in benchmarking is the fact that many researchers actually do a compromise between quality and speed when they evaluate the accuracy and runtime of their approach. Often the design of efficient algorithms is accompanied by the proposal of improved model components. In this case, some of the quality gain is sacrificed for improving the run time. Typically iterative schemes are stopped much too early and the results are far from convergence. In such cases it is not clear how good the novel model actually performs and what part of the speed up refers to the underlying (potential parallel) implementation. It is evident that this problem is a general problem that is also relevant for sequential algorithms. In this context, the members of the discussion groups proposed that every research paper benchmarking efficient algorithms should provide at least one experiment where the true potential of an algorithm is made clear. This fully converged high quality result may then serve as a reference solution to judge the loss of quality when cutting down the runtime of the algorithm. Evidently, in the ideal case, no such loss should occur. However, one may observe that for some tasks approximate solutions obtained either by approximating the original problem or by stopping the algorithm too early are absolutely fine. In those cases it may even seem possible to obtain further speed ups by reducing the complexity of the original model. This, however, is another open question that came up during discussion and requires further research.

Participants

- Egil Bae
University of Bergen, NO
- Yuri Boykov
Middlesex College – Ontario, CA
- Kristian Bredies
Universität Graz, AT
- Xavier Bresson
City Univ. – Hong Kong, HK
- Thomas Brox
Universität Freiburg, DE
- Andrés Bruhn
Universität des Saarlandes, DE
- Antonin Chambolle
Ecole Polytechnique – Palaiseau, FR
- Raymond Chan
Chinese Univ. of Hong Kong, HK
- Christian Clason
Universität Graz, AT
- Daniel Cremers
TU München, DE
- Bastian Goldlücke
TU München, DE
- Sven Grewenig
Universität des Saarlandes, DE
- Michael Hintermueller
HU Berlin, DE
- Sung Ha Kang
Georgia Inst. of Technology, US
- Ron Kimmel
Technion – Haifa, IL
- Harald Köstler
Univ. Erlangen-Nürnberg, DE
- Vladimir Kolmogorov
IST Austria – Klosterneuburg, AT
- Arjan Kuijper
Fraunhofer Inst. – Darmstadt, DE
- Christoph Lampert
IST Austria – Klosterneuburg, AT
- Stamatis Lefkimmiatis
EPFL – Lausanne, CH
- Jan Lellmann
Universität Heidelberg, DE
- Stacey Levine
Duquesne University, US
- Cécile Louchet
Université d'Orléans, FR
- Arvid Lundervold
University of Bergen, NO
- Jan Modersitzki
Universität zu Lübeck, DE
- Thomas Pock
TU Graz, AT
- Guy Rosman
Technion – Haifa, IL
- Ulrich Rüde
Univ. Erlangen-Nürnberg, DE
- Gabriele Steidl
TU Kaiserslautern, DE
- Evgeny Strekalovskiy
TU München, DE
- Robert Strzodka
MPI für Informatik – Saarbrücken, DE
- Xue-Cheng Tai
University of Bergen, NO
- Wenbing Tao
Huazhong University of Science & Technology, CN
- Tanja Teuber
TU Kaiserslautern, DE
- Olga Veksler
Univ. of Western Ontario, CA
- Andreas Wedel
Daimler AG, DE
- Joachim Weickert
Universität des Saarlandes, DE
- Manuel Werlberger
TU Graz, AT
- Wotao Yin
Rice University – Houston, US