Report from Dagstuhl Seminar 14072

New Perspectives in Shape Analysis

Edited by

Michael Breuß[1], Alfred M. Bruckstein[2], Petros Maragos[3], and Stefanie Wuhrer[4]

1 **BTU Cottbus, DE,** breuss@tu-cottbus.de
2 **Technion – Haifa, IL,** freddy@cs.technion.ac.il
3 **National TU – Athens, GR,** maragos@cs.ntua.gr
4 **Universität des Saarlandes, DE,** swuhrer@mmci.uni-saarland.de

Abstract

Over the last decade, it has become increasingly affordable to digitize 2-D and 3-D shape information using multiple modalities, such as (video) cameras, image-based reconstruction systems, laser-range scanners, or depth cameras. If these dense models can be processed and described in an efficient and informative way, they can be used in applications, such as ergonomic design, virtual shopping, scientific and medical visualization, realistic simulation, photo-realistic rendering, the design of natural user interfaces, and semantic scene understanding.

Traditionally, the notion of shape has been studied either by analyzing projections of shapes in images or by analyzing a sparse set of marker positions on 3-D shapes. Typical tasks in 2-D shape analysis include segmenting objects in images and tracking objects across a sequence of images, and typical tasks in 3-D shape analysis include reconstructing the three-dimensional object depth from input images and identifying corresponding points on different 3-D models. The analysis and processing of shape data becomes especially challenging because of the increasing amount of data captured by sensors used to acquire shapes, and because modern applications such as natural user interfaces require real-time processing of the input shapes.

Meeting these challenges requires models of shape analysis that are compact and informative, thereby allowing the development of algorithms that can process large datasets efficiently. To achieve these goals, interdisciplinary approaches are needed that use concepts from a variety of research areas, including numerical computing, differential geometry, deformable shape modeling, sparse data representation, and machine learning. On the algorithmic side, many shape analysis tasks are modeled using partial differential equations, which can be solved using tools from the field of numerical computing. The fields of differential geometry and deformable shape modeling have recently begun to influence shape analysis methods. Furthermore, tools from the field of sparse representations, which aim to describe input data using a compressible representation with respect to a set of carefully selected basis elements, have the potential to significantly reduce the amount of data that needs to be processed in shape analysis tasks. The related field of machine learning offers similar potential.

This seminar brought together 28 researchers from North America and Europe engaged in recent and upcoming developments in shape analysis who view these challenges from different perspectives and who together discussed the pressing open problems and novel solutions to them.

Seminar February 9–14, 2014 – http://www.dagstuhl.de/14072

1998 ACM Subject Classification G.1.3 Numerical Linear Algebra, G.1.8 Partial Differential Equations, I.2.10 Vision and Scene Understanding, I.3.5 Computational Geometry and Object Modeling, I.3.7 3-D Graphics and Realism, I.4.6 Segmentation, I.4.8 Scene Analysis, I.4.10 Image Representation, I.5.4 Applications

Keywords and phrases shape analysis, mathematical morphology, shape reconstruction, sparsity, machine learning, numerical computing, level set methods, optimisation methods

Digital Object Identifier 10.4230/DagRep.4.2.60

Edited in cooperation with Yong Chul Ju

New Perspectives in Shape Analysis, *Dagstuhl Reports*, Vol. 4, Issue 2, pp. 60–78
Editors: Michael Breuß, Alfred M. Bruckstein, Petros Maragos, and Stefanie Wuhrer
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Michael Breuß
Alfred M. Bruckstein
Petros Maragos
Stefanie Wuhrer

Dagstuhl seminar 14072 *New Perspectives in Shape Analysis* took place February 9–14, 2014. 28 researchers from North America and Europe discussed state-of-the-art, current challenges, and promising future research directions in the areas of 2-D and 3-D shape analysis from a cross-disciplinary point of view. Participants included international experts from the fields of continuous-scale shape analysis, discrete shape analysis and sparsity, and numerical computing. The seminar consisted of an opening session, 11 scientific presentation sessions, as well as a break-out session, which provided room for in-depth discussions in small groups. Furthermore, there was time for extensive discussions both between the talks and in the evenings.

This seminar was motivated by the observation that in everyday life, geometric shapes surround us, and that the understanding of concepts describing these shapes is at the heart of various applications, such as ergonomic design, virtual shopping, scientific and medical visualization, realistic simulation, the design of natural user interfaces, and semantic scene understanding. Traditionally, the notion of shape has been studied either by analysing a sparse set of marker positions on 3-D shapes, mostly for medical imaging applications, or by analysing projections of shapes in 2-D images, mostly for image processing and computer vision applications. New challenges in the analysis and processing of such data arise with the increasing amount of data captured by sensors used to acquire shapes, and with modern applications such as natural user interfaces that require real-time processing of the input shapes. Recently, it has become increasingly affordable to digitise 3-D shapes using multiple modalities, such as laser-range scanners, image-based reconstruction systems, or depth cameras like the Kinect sensor. Using these dense 3-D shapes in the above mentioned applications requires processing and describing the shapes in an efficient and informative way.

The purpose of this seminar was to address these challenges with the latest tools related to geometric, algorithmic and numerical concepts. To do so, we brought together researchers working on shape analysis topics from different perspectives.

As the analysis of 3-D shapes and deformable shape models have received much interest recently, classic shape analysis tools from differential geometry have a fresh influence in the field. Being related to the issue how to represent shapes efficiently, the research areas of sparse data representation and machine learning have begun to influence shape analysis modelling and the numerics. Especially in the context of three-dimensional data (or even higher-dimensional data sets), efficient optimization methods will certainly become increasingly important since many shape analysis models can be cast in the form of an optimization problem.

While the fields of modelling and numerical computing are strongly related when it comes to shape analysis applications, modelling is seen as a hot topic in computer science while numerical computing is often seen as a mathematical domain. The purpose in bringing together researchers from those different communities sharing substantial interest in shape analysis was to explore the benefits of a *cross-disciplinary* point of view. More specifically,

- researchers in continuous-scale shape analysis brought to the meeting their knowledge of differential and variational models and also of classic numerical methods in the field,
- researchers in discrete shape analysis and sparsity brought to the meeting their knowledge about the latest techniques in efficient data representations and related machine learning techniques, as well as efficient data structures and discrete optimization methods, and
- researchers in numerical computing brought to the meeting their knowledge of numerical techniques for PDEs and optimization.

As the demands in the individual fields are high, the research groups in which the most interesting techniques are proposed are quite specialised. This not only holds for discrete and continuous-scale modelling and numerical computing, but also for the areas of sparsity and machine learning that were discussed during this seminar. Because of this, there is no regular conference or workshop that serves as a meeting place for an exchange of ideas of these groups.

Promising new ways to combine the latest techniques from these different fields were identified during in-depth discussions in small groups. Some especially promising research directions in the areas of intrinsic structure detection, co-segmentation of shapes, shape from shading, modelling deformable shapes, and models for face shapes, were discussed in small groups during the break-out session.

2 Table of Contents

3 Overview of Talks

3.1 An operator approach to geometry processing

Mirela Ben-Chen (Technion – Haifa, IL)

Joint work of Ben-Chen, Mirela; Azencot, Omri; Ovsjanikov, Maks; Chazal Fred
Main reference O. Azencot, M. Ben-Chen, F. Chazal, M. Ovsjanikov, "An Operator Approach to Tangent Vector Field Processing," Computer Graphics Forum, 32(5):73–82, 2013.
URL http://dx.doi.org/10.1111/cgf.12174

A fundamental task in the analysis of geometric shapes is the choice of representation of differential quantities. Different representations will lead to different optimization problems, and can considerably influence the scope of feasible applications. We will describe a novel choice of representation, based on functional operators, which although standard in classic differential geometry has only been recently applied to geometry processing. We will show how operators which take scalar functions to scalar functions can be used to concisely represent smooth maps between shapes and smooth vector fields. We further demonstrate that by using a common operator representation, the intimate connection between maps and vector fields can be leveraged to easily compute vector fields which fulfill intricate global constraints, such as Killing and symmetric vector fields. Finally, we discuss additional geometric properties which can be represented as operators.

3.2 Combination of piecewise-geodesic paths for interactive segmentation

Sébastien Bougleux (Caen University, FR)

Joint work of Mille, Julien; Bougleux, Sebastien; Cohen, Laurent D.
Main reference J. Mille, S. Bougleux, L. D. Cohen, "Combination of paths for interactive segmentation," in Proc. of the British Machine Vision Conf. 2013, pp. 133.1–133.11, BMVA Press, 2013.
URL http://dx.doi.org/10.5244/C.27.133

When it is desired to segment interactively an object of interest in an image, it is natural to position points manually along the contour of the object. By constructing a path between each pair of consecutive control points and by joining these paths end-to-end, a closed contour is obtained. The resulting contour may not be simple and may not follow the object's contour everywhere. These problems may occur when control points are irregularly distributed or when the contour is poorly represented. To overcome these problems, we propose a method guaranteeing a simple and closed contour while being robust to the initial placement of the control points. The central idea is to consider several possible paths per pair of consecutive control points, and to select the combination of paths so that the resulting contour is closed, and as simple as possible. To this, we focus on a wider category of paths, called *piecewise geodesic*, and we propose an energy functional combining classical *edge* and *region* terms with a new term which favors curves to be simple. This term penalizes recoveries and self-intersections. Given n control points and K paths per couple of consecutive points, we also introduce a heuristic algorithm in $O(Kn^2)$ for determining a satisfactory combination, instead of an exhaustive search in $O(K^n)$. Comparison against the geodesically linked active contour model and the Riverbed algorithm, which have similar inputs and purposes, demonstrated the advantages of our approach.

3.3 Is there a good coordinate system for shape from shading?

Michael Breuß (BTU Cottbus, DE)

Joint work of Galliani, Silvano; Ju, Yong Chul; Breuß, Michael; Bruhn, Andrés
Main reference S. Galliani, Y. C. Ju, M. Breuß, A. Bruhn, "Generalised Perspective Shape from Shading in Spherical Coordinates," in Proc. of the 4th Int'l Conf. on Scale Space and Variational Methods in Computer Vision (SSVM'13), LNCS, Vol. 7893, pp. 222-233, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-38267-3_19

In this talk some recent results are presented making use of spherical coordinates for perspective shape from shading. The aim of the talk is to inspire some discussion about good coordinate systems for shape from shading, and probably for shape-from-X-tasks more generally.

3.4 Joint diagonalization methods for deformable shape analysis

Michael M. Bronstein (University of Lugano, CH)

Joint work of Bronstein, Michael M.; Bronstein, Alexander M.; Kovnatsky, Artiom; Kimmel, Ron; Glashoff, Klaus
Main reference A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, R. Kimmel, "Coupled quasi-harmonic bases," Computer Graphics Forum, 32(2pt4):439–448, 2013.
URL http://dx.doi.org/10.1111/cgf.12064

I will present our recent work in which we introduced joint diagonalization methods, previously used for blind separation of signals, into the analysis of deformable shapes. I will show how these methods can be applied to problems of shape correspondence, similarity, and editing, and discuss their relations to other interesting recent works in the field.

3.5 Hierarchical Representation of Cell Complexes for Topological Shape Analysis

Leila De Floriani (University of Genova, IT)

Joint work of De Floriani, Leila; Lidija, Comic; Federico Iuricich; Ulderico Fugacci
Main reference L. Comic, L. De Floriani, F. Iuricich, U. Fugacci, "Topological Modifications and Hierarchical Representation of Cell Complexes in Arbitrary Dimensions," Computer Vision and Image Understanding, Vol. 121, pp. 2–12, 2014.
URL http://dx.doi.org/10.1016/j.cviu.2013.11.011

In this talk, we present a hierarchical approach to homology computation and topological analysis on cell complexes of arbitrary dimensions. Specifically, we briefly review a set of atomic Euler operators for simplifying and refining cell complexes. We have shown that such operators form a minimally complete set for updating cell complexes. Based on them, we have developed a hierarchical model for cell complexes, the we call a Hierarchical Cell Complex (HCC), which allows extracting representations of the original complex at both uniform and variable resolutions. We discuss our implementation of the HCC based on homology-preserving Euler operators, and we show the effectiveness of such a hierarchical model in computing homology generators efficiently and at different resolutions. We discuss on-going work on extending the HCC to the case of simplicial complexes.

3.6 Image segmentation with level set trees

Anastasia Dubrovina (Technion – Haifa, IL)

Joint work of Dubrovina, Anastasia; Hershkovitz, Rom; Kimmel, Ron

Several image editing tasks, such as segmentation, matting and colorization, employ an intrinsic distance measure defined in the image domain, to propagate information from user defined set of constraints, to the whole image. In this talk we will show that this problem can be alternatively defined in the domain of the level sets of an image, and solved by traversing the tree of level sets. We will further show application of proposed method for user-assited image segmentation. Based on joint work with Ron Kimmel and Rom Hershkovitz.

3.7 Are Fast Solvers for Hamilton-Jacobi equations really reliable?

Maurizio Falcone (University of Rome "La Sapienza", IT)

Joint work of Cacace, Simone; Cristiani, Emiliano; Falcone, Maurizio
Main reference S. Cacace, E. Cristiani, M. Falcone, "Can local single-pass methods solve any stationary Hamilton-Jacobi-Bellman equation?" to appear in SIAM J. Sci. Comput.; availabe as pre-print as arXiv:1301.6775v3 [math.NA].
URL http://arxiv.org/abs/1301.6775v3

The use of local single-pass methods (like, e.g., the Fast Marching method) has become popular in the solution of some Hamilton-Jacobi equations. The prototype of these equations is the eikonal equation related to the level-set method and applied in several applications including image processing. For this equation the methods can be applied saving CPU time and possibly memory allocation. Then, some natural questions arise: can local single-pass methods solve any Hamilton-Jacobi equation? If not, where the limit should be set? We will try to answer these questions. In order to give a complete picture, we present an overview of some fast methods available in literature and we briefly analyze their main features. We also provide some numerical tests which are intended to exhibit the limitations of the methods. We show that the construction of a local single-pass method for general Hamilton-Jacobi equations is very hard, if not impossible. Nevertheless, some special classes of problems can be actually solved, making local single-pass methods very useful from the practical point of view. Some examples taken from Shape-from-Shading models will be discussed.

References

1 Cacace, S., Cristiani, E., Falcone, M.: Can local single-pass methods solve any stationary Hamilton-Jacobi-Bellman equation? SIAM J. Sci. Comput., to appear
2 Cacace, S., Cristiani, E., Falcone, M.: Two Semi-Lagrangian Fast Methods for Hamilton-Jacobi-Bellman Equations,to appear in Proceedings IFIP 2013

3.8 Image-Based Approaches for Photo-Realistic Rendering of Clothes

Anna Hilsmann (HU Berlin, DE)

Joint work of Hilsmann, Anna; Fechteler, Philipp; Eisert, Peter
Main reference A. Hilsmann, P.Fechteler, P. Eisert, "Pose Space Image Based Rendering," Computer Graphics Forum, 32(2):265–274, 2013.
URL http://dx.doi.org/10.1111/cgf.12046

In this talk, I presented methods for image-based rendering and modification of objects with complex appearance properties, concentrating on the example of clothes. With physical simulation methods, rendering of clothes is computationally demanding because of complex cloth drapery and shading. In contrast, the proposed methods use real images, which capture these properties and serve as appearance examples to guide complex animation or texture modification processes. Texture deformation and shading are extracted as image warps both in the spatial and in the intensity domain (Hilsmann et al 2010). Based on these warps, a pose-dependent image-based rendering method synthesizes new images of clothing from a database of pre-recorded images (Hilsmann et al. 2013). For rendering, the images and warps are parametrized and interpolated inpose-space, i.e. the space of body poses, using scattered data interpolation. To allow for appearance changes, an image-based retexturing method is proposed,which exchanges the cloth texture in an image while maintaining texture deformation and shading properties, without the knowledge of the scene geometry and lighting conditions (Hilsmann et al. 2011). Altogether, the presented approaches shift computational complexity from the rendering to an a-priori training phase. The use of real images and warp-based extraction of deformation and shading allow a photo-realistic visualization and modification of clothes, including fine details, without computationally demanding simulation of the underlying scene and object properties.

3.9 Variational Perspective Shape from Shading

Yong Chul Ju (Universität Stuttgart, DE)

Joint work of Breuß, Michael; Bruhn, Andrés

Shape from Shading (SfS) has been a classical task over decades in computer vision and enormous progress has been made with the methods of using partial differential equations in recent years. However, there is still a lack of state-of-the-art models in variational methods. In this talk, we propose a novel variational model for SfS which offers several advantages compared to existing variational approaches.

(i) It is the first variational model based on advanced camera models such as perspective projections.
(ii) As a perspective model it is the first one that directly allows to compute depth instead of estimating and integrating surface normals.
(iii) Due to the direct computation it does not require integrability constraints since it fulfils the integrability condition by construction.

This is joint work with Michael Breuß (Brandenburg University of Technology) and Andrés Bruhn (University of Stuttgart).

3.10 Segmentation in 3-D surface data

Margret Keuper (Universität Freiburg, DE)

Currently, several methods are available to reconstruct 3-D surfaces of complex scenes with a certain accuracy. However, segmentation in these scenes is not trivial. First, the data is usually not complete, i.e. we are not provided aclosed triangle mesh but rather a triangle soup with holes. Secondly, the datais degraded by recording noise resulting in reconstruction artefacts. Thirdly, the lack of fully labelled 3-D data makes it hard to learn object category specific appearance. I am going to present our most recent work on graph reduction for image segmentation and its extension to point clouds from single view RGB-D. Our method allows for a fast and accurate segmentation of large point clouds using the spectral clustering paradigm.

3.11 Data flattening: A spectral perspective

Ron Kimmel (Technion – Haifa, IL)

Joint work of Aflalo, Yonathan; Kimmel, Ron
Main reference Y. Aflalo, R. Kimmel, "Spectral multidimensional scaling," in Proc. of the National Academy of Sciences of the United States of America (PNAS), 110(45):18052–18057, 2013.
URL http://dx.doi.org/10.1073/pnas.1308708110

Mapping metric structures to the plane can be realized by multidimensional scaling methods. The input to such procedures is the set of inter-geodesic distances between points of the data we would like to map to the plane. The complexity of the input to these algorithms is quadratic in the number of inputdata points. In the talk a quasi-linear method (in space and time) was presented for that goal exploiting properties of the natural basis for this problem. Based on joint work with Yonathan Aflalo.

References

1 Y. Aflalo, and R. Kimmel. Spectral multidimensional scaling, Proc. NationalAcademy of Sciences (PNAS), 110(45):18052-18057, 2013

3.12 Efficient Detection of Deformable Objects

Iasonas Kokkinos (Ecole Centrale – Paris, FR)

Joint work of Kokkinos, Iasonas; Haithem Boussaid
Main reference I. Kokkinos, "Rapid Deformable Object Detection using Dual Tree Branch and Bound," in Proc. of the 25th Annual Conf. on Neural Information Processing Systems (NIPS'11), pp. 2681–2689, 2011.
URL http://papers.nips.cc/paper/4338-rapid-deformable-object-detection-using-dual-tree-branch-and-bound.pdf

Deformable Part Models (DPMs) play a prominent role in current object recognition research, as they allow to rigorously deal with shape variability in an object category. Still, the computational requirements of such models make them unsuitable for real-time applications.

In this talk we will see how well-known optimization techniques, such as Branch-and-Bound and the Alternating Direction Method of Multipliers (ADMM) can be used to efficiently perform inference with DPMs. Instead of evaluating the classifier score exhaustively for all part locations and scales, such techniques allow us to quickly focus on promising image locations. The core problem that we will address is how to compute bounds that accommodate part deformations; this allows us to apply Branch-and-Bound to our problem. We will then present a technique to provide quick probabilistic bounds to the individual part scores by using Chebyshev's inequality, and will finally present a decomposition technique that allows us to deal with models of complicated graph structure by employing ADMM for fast convergence. When comparing to the baseline DPM implementation of Felzenszwalb et al, we obtain exactly the same results but can perform the part combination substantially faster (yielding double up to tenfold or higher speedups); on a challenging medical shape segmentation benchmark our ADMM-based technique yields substantially better results than the previous state-of-the-art, while converging in a few seconds.

3.13 On Shape Recognition and Language

Petros Maragos (National TU – Athens, GR)

URL http://cvsp.cs.ntua.gr

Shapes convey meaning. Language is efficient in expressing and structuring meaning. The main thesis of this presentation is that by fusing shape and linguistic information shape recognition can be improved in performance and can also be enriched by adding language-based semantics. This talk focuses on two paradigms from two broad classes of moving shapes that are related to language: (1) moving handshapes from videos of sign language, and (2) multimodal gestures combining audio and visual information. In (1) we fuse shape information that consists of 2-D handshape images and their 2-D motion-position with sign linguistic information in the form of subunit sequences. In (2) we fuse 3-D shape from motion-position of gesturing hands/arms and 2-D handshapes in RGB and Depth visual channels with audio information in the form of acoustically recognized keyword sequences. In both applications the performance of shape classification or recognition significantly improves by fusing geometry with linguistic information. This is joint work with S. Theodorakis, V. Pitsikalis, A. Roussos, G. Pavlakos and A. Katsamanis. Supported by the research project "COGNIMUSE". More information, related papers and current results can be found in http://cvsp.cs.ntua.gr.

3.14 Relation between total variation and persistence distance and its application in signal processing

Gerlind Plonka-Hoch (Universität Göttingen, DE)

Joint work of Plonka-Hoch, Gerlind; Zheng, Yi
Main reference G. Plonka, Y. Zheng, "Relation between total variation and persistence distance and its application in signal processing," Manuscript, January 2014.
URL http://num.math.uni-goettingen.de/~plonka/pdfs/plonka_zheng.pdf

In this talk we establish the new notion of persistence distance for discrete signals and study its main properties. The idea of persistence distance is based on recent developments in topological persistence for assessment and simplification of topological features of data sets. Particularly, we establish a close relationship between persistence distance and discrete total variation for finite signals. This relationship allows us to propose a new adaptive denoising method based on persistence that can also be regarded as a nonlinear weighted ROF model. Numerical experiments illustrate the ability of the new persistence based denoising method to preserve significant extrema of the original signal.

3.15 Time Discrete Geodesic Paths in the Space of Images

Martin Rumpf (Universität Bonn, DE)

Joint work of Berkels, Benjamin; Effland, Alexander; Rumpf, Martin

The space of images is considered as a Riemannian manifold using the metamorphosis approach proposed by Trouvé and Younes. Via a generalization of the flow of diffeomorphism approach simultaneous transport and intensity variations are both reflected by the underlying Riemannian metric. A robust and effective variational time discretization of geodesics paths is proposed, which requires to minimize a discrete path energy consisting of a sum of consecutive image matching functionals over a set of image intensity maps and pairwise matching deformations. Under minimal regularity requirements for the input images the existence of discrete geodesic paths defined as minimizers of this variational problem is shown. Furthermore, Γ-convergence of the underlying discrete path energy to the continuous path energy is proved. This includes a diffeomorphism property for the induced transport and the existence of a square-integrable weak material derivative in space and time. A spatial discretization via finite elements combined with an alternating descent scheme in the set of image intensity maps and the set of matching deformations is presented to numerically approximate the discrete geodesic paths. Computational results underline the efficiency of the proposed approach and demonstrate important qualitative properties.

3.16 Shape distance and Shape prior

Frank R. Schmidt (Universität Freiburg, DE)

Joint work of Gorelick, Lena; Schmidt, Frank R.; Boykov, Yuri
Main reference L. Gorelick, F. R. Schmidt, Y. Boykov, "Fast Trust Region for Segmentation," in Proc. of the 2013 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'13), pp. 1714–1721, IEEE, 2013.
URL http://dx.doi.org/10.1109/CVPR.2013.224

Trust region is a well-known general approach to optimization which offers many advantages over standard gradient descent techniques. In particular, it allows more accurate nonlinear approximation models. In each iteration this approach computes a global optimum of a suitable approximation model within a fixed radius around the current solution, a.k.a. trust region. In general, this approach can be used only when some efficient constrained optimization algorithm is available for the selected non-linear (more accurate) approximation model. In this talk I present a Fast Trust Region (FTR) approach for optimization of segmentation energies with non-linear regional terms, which are known to be challenging for existing algorithms. These energies include, but are not limited to, volume constraint and Bhattacharyya distance between the observed and the target appearance distributions. In particular, I show that replacing the well-known Hamming distance with a distance that is more related to shape matching results in a robust algorithm that outperforms the classical level-set approach with respect to both accuracy and time.

3.17 Anisotropic Third-Order Regularization for Sparse Digital Elevation Maps

Carola-Bibiane Schönlieb (University of Cambridge, GB)

Joint work of Lellmann, Jan; Morel, Jean-Michel; Schönlieb, Carola-Bibiane
Main reference J. Lellmann, J.-M. Morel, C.-B. Schönlieb, "Anisotropic Third-Order Regularization for Sparse Digital Elevation Models," in Proc. of the 4th Int'l Conf. on Scale Space and Variational Methods in Computer Vision (SSVM'13), LNCS, Vol. 7893, pp. 161–173, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-38267-3_14

We consider the problem of interpolating a surface based on sparse data such as individual points or level lines. We derive interpolators satisfying a list of desirable properties with an emphasis on preserving the geometry and characteristic features of the contours while ensuring smoothness across level lines. We propose an anisotropic third-order model and an efficient method to adaptively estimate both the surface and the anisotropy. Our experiments show that the approach outperforms AMLE and higher-order total variation methods qualitatively and quantitatively on real-world digital elevation data.

3.18 3-D statistical face models for 2-D image analysis

William Smith (University of York, GB)

Main reference O. Aldrian, W. A. P. Smith, "Inverse Rendering of Faces with a 3-D Morphable Model," IEEE Trans. on Pattern Analysis and Machine Intelligence, 35(5):1080–1093, 2013.
URL http://dx.doi.org/10.1109/TPAMI.2012.206

Statistical face models have a long history in computer vision, including the now classical techniques of Eigenfaces, Active Shape Models and Active Appearance Models. Later, approaches such as 3-D Morphable Models attempted to capture the intrinsic features of a face (3-D shape and reflectance properties) allowing an explicit image formation model to be used in an analysis-by-synthesis framework. In this talk, I will describe recent work on obtaining training data (i.e. capturing face intrinsics) and how such models can be used to analyse 2-D images, providing a statistical route to facial shape-from-shading. Finally, I will discuss what I believe are the open questions in statistical face modelling.

3.19 Towards a design grammar

Sibel Tari (Middle East Technical University – Ankara, TR)

Joint work of Tari, Sibel; Ozkar, M.; H. Keles
Main reference H. Y. Keles, M. Özkar, S. Tari, "Weighted shapes for embedding perceived wholes," Environment and Planning B: Planning and Design, 39(2):360–375, 2012.
URL http://dx.doi.org/10.1068/b37067

A major step towards computer-aided creative design is the ability to automatically envision different parts in a whole and to continuously divide and combine. In this talk, I will review recent progress at Middle East Technical University towards a design grammar implementation. The talk will feature a couple of interrelated methods, all benefiting from the use of diffuse forms of designs computed via Screened Poisson. I will also address how new designs can be generated from older ones using mild randomization. (The talk is based on two joint works, with M. Ozkar and with H. Keles)

3.20 Semi-Lagrangian Approximation of non-Lambertian Reflectance Models for the Orthographic Shape-from-Shading Problem

Silvia Tozza (University of Rome "La Sapienza", IT)

Several advances have been made in the last ten years to improve the Shape–from–Shading model in order to allow its use on real images. The classic Lambertian model, suitable to reconstruct 3-D surfaces with uniform reflection properties, has shown to be unsuitable for other types of surfaces, for example for rough objects consisting of materials such as clay. Other models have been proposed by several authors but it is still unclear what could be the best one (if there is a best). To this end we have analyzed some models for non-Lambertian surfaces focussing our attention on Phong and Oren-Nayar models which seem to be more flexible and accurate than other models. In particular, the first is intended to treat specular surfaces whereas the second includes diffusion reflectance terms. We present a semi-Lagrangian approximation schemes for nonlinear partial differential equations corresponding to these models, analyze their properties and compare them with the classical Lambertian model on a series of benchmarks images.

3.21 Dense Non-Rigid Shape Correspondence using Random Forests

Matthias Vestner (TU München, DE)

Since the 2000s there has been an increasing influx of work on finding and describing correspondences among 3-dimensional shapes. Due to the highly ill-posed nature of the shape matching problem, it is very unlikely that a general method will reliably find good matchings between arbitrary shapes. In this talk the correspondence problem between non-rigid shapes will be tackledvia a learning-by-examples approach. We treat the shape matching problem as a classification problem, where input samples are points on a test-shape and the output class is an element of a canonical label set, which might e.g. coincide with one of the shapes in the training set. As a classifier we choose the random forest paradigm, that can in general be built upon any parametrizable feature descriptor. As a particular instance of such a descriptor the Wave Kernel signature will be presented. In order to get rid of artifacts caused by symmetry properties of the considered shapes further a regularizer will be introduced.

3.22 Morphological amoebas and image segmentation

Martin Welk (UMIT – Hall in Tirol, AT)

Morphological amoebas are a class of image-adaptive structuring elements which can be used in connection with various morphological filters. Amoeba-based filters have interesting connections to partial differential equations used in image processing; for example, amoeba median filtering is closely related to the self-snakes PDE [6, 5, 2]. Based on this observation, also an active contour algorithm can be devised that uses amoeba median filtering to evolve the level set function of a contour [5, 4, 3].

In the talk results on this segmentation approach were summarised. Additionally, preliminary results of recent work directed at extracting additional information on texture from the intrinsic structure of amoebas were presented [1].

References

1 M. Welk. Discrimination of image textures using graph indices. In M. Dehmer, F. Emmert-Streib, eds., Quantitative Graph Theory: Mathematical Foundations and Applications, CRC Press, in press.
2 M. Welk, M. Breuß. Morphological amoebas and PDEs. In P. Hawkes, ed., Advances in Imaging and Electron Physics, vol. 185, in press.

3 M. Welk. Analysis of Amoeba Active Contours. Technical Report, http://arXiv.org/abs/1310.0097
4 M. Welk. Relations between amoeba median algorithms and curvature-based PDEs. In A. Kuijper et al., eds., Scale-Space and Variational Methods in Computer Vision, LNCS 7893, 392–403, Springer, Berlin, 2013.
5 M. Welk. Amoeba active contours. In A. Bruckstein et al., eds., Scale Space and Variational Methods in Computer Vision. LNCS 6667, 374–385, Springer, Berlin, 2012.
6 M. Welk, M. Breuß, O. Vogel. Morphological amoebas are self-snakes. Journal of Mathematical Imaging and Vision, 39:87–99, 2011.

3.23 Statistical Modeling of 3-D Human Face Shapes

Stefanie Wuhrer (Universität des Saarlandes, DE)

Joint work of Brunton, Alan; Bolkart, Timo; Wuhrer, Stefanie
Main reference A. Brunton, T. Bolkart, S. Wuhrer, "Multilinear Wavelets: A Statistical Shape Space for Human Faces," arXiv:1401.2818v1 [cs.CV], 2014.
URL http://arxiv.org/abs/1401.2818v1

Accurate reconstruction of face shape is important for applications such as tele-presence and gaming. Such a reconstruction problem can be solved efficiently and in the presence of noise with the help of statistical shapemodels that constrain the shape of the reconstruction. In this talk, a statistical model to represent 3-D human faces in varying expression is discussed, which decomposes the surface of the face using a wavelet transform, and learns many localized, decorrelated multilinear models on the resulting coefficients. The localized and multi-scale nature of this model allows for recovery of fine-scale detail while retaining robustness to severe noise and occlusion, and is computationally efficient and scalable.

3.24 Shape Compaction

Hao Zhang (Simon Fraser University – Burnaby, CA)

Joint work of Zhang, Hao; Hu, Ruizhen; Li, Honghua Li; Alhasheem, Ibraheem; Shamir, Ariel; Cohen-Or, Daniel
Main reference H. Li, I. Alhashim, H. Zhang, A. Shamir, D. Cohen-Or, "Stackabilization," ACM Trans. on Graphics (Special Issue of SIGGRAPH Asia), 31(6), Article 158, 2012.
URL http://doi.acm.org/10.1145/2366145.2366177
URL http://www.computer-graphics.cn/~hh/projects/stackem/

Compact shape *representations*, e.g., compression, is a well-studied problem. We are interested in compaction at the shape *configuration* level. The termshape configuration refers to how a shape, real or conceptual, is physically modelled (e.g., in terms of design and composition of its parts) and spatially arranged (e.g., in terms of shape/part positioning and possibly in relation to other shapes). The required storage is the actual physical space the shape configuration occupies. Compact shape configurations can save valuable space in industrial settings, e.g., for storage, shipping, printing, etc., leading to cost reduction. The key difference to compaction at the representation level is that by changing shape configurations, we allow a given shape to undergo significant changes, which are certainly beyond a close geometric approximation. For example, we may change how a shape is decomposed and assembled so it can be better folded or alter its geometry so that it can be more compactly stacked with other shapes. Compact shape configurations can be generated by either re-modeling or re-arranging the parts within one shape or changing the inter-shape spacial relations of a set shapes. In this talk, I will pose three new problems: stackabilization, pyramidaliztion, and foldablization, and present our progress on solving them.

3.25 Shape descriptors/measures

Jovisa Zunic (University of Exeter, GB)

Shape is an object property which has a big discrimination capacity. This is because the shape allows many numerical characterisations. Some of these numerical characterisations relate to the validation of certain intuitively clear shape properties (usually called shape descriptors) e.g. shape convexity, shape elongation, shape compactness, etc. Methods for numerical validations, of such properties (i.e. shape descriptors) are called the shapemeasures. Some of shape descriptors have already assigned multiple shapemeasures (convexity, compactness, etc). Since shape descriptors have an intuitively clear meaning, the behaviour of the measures, assigned to them, can be predicted to some extent, which enables a priory estimate of how a certain measure fits to the task performed. The problems is that the number of shape descriptors is not large, e.g. comparing with the number of generic shape descriptors (usually do not have a clear geometric interpretation). In this talk, some ideas/results about how to overcome such a drawback and how to expand the applicability of shape/descriptors will be presented.

4 Working Groups

4.1 Summary of Break-out Session

Yong Chul Ju (Universität Stuttgart, DE)
Stefanie Wuhrer (Universität des Saarlandes, DE)

During the break-out session, five small groups discussed the following topics with the goal to identify the most pressing research questions in these areas: intrinsic structure, (co)segmentation, shape from shading, models for deformable shapes, and face deformation models. Here, we briefly summarise the outcomes of the discussions.

Intrinsic Structure: Efficient shape representations play an important role for many tasks in computer vision, such as the matching problem or object recognition. However, simple geometric features of an object, in general, are not always distinctive enough to capture important structures, such as the distortions of non-rigid objects. The study of the self-similarity of an object under deformations that preserve either the extrinsic or the intrinsic geometry of the object may help to derive efficient shape representations. This observation leads to the main question of how to encode the self-similarity of an object as shape descriptor.

(Co)Segmentation: Co-segmentation is the task of segmenting a set of given images or 3-D shapes into corresponding parts. Since it is extremely difficult to have a perfect match of an object between one or more pairs of input images, e.g. due to self-occlusion and shading, one has to almost always deal with missing correspondences. As a consequence, incorporating partial correspondences is an essential step to tackle the (co)segmentation problem from several input images successfully.

Shape from Shading (SfS): Although SfS has been one of the key problems in computer vision, there is still no unified SfS model framework that can combine several modelling components and parameters. Specifically, it remains a challenge to combine non-Lambertian and Lambertian models and to formulate one equation composed of ambient, Lambertian and specular parts. In order to further develop new models and solutions, a carefully devised unified framework would be a valuable tool. In addition, most work related to SfS considers only a single greyscale input image. However, exploiting colour information for SfS would be interesting, as most greyscale images are also acquired from colour images nowadays.

Models for deformable shapes: In order to model deformable shapes, mimicking the behaviour of an amoeba, a unicellular organism, is a basic and effective tool. However, there is still no satisfying mathematical description of how amoebas move in the plane. So finding such a description is the first step toward further development in the field. In addition, interpenetration effects in real world situations caused by interweaving or distortions, e.g. in the overlapping area on the human face, remain extremely challenging. There is a definite need to handle these situations.

Face deformation models: As a special case of deformable shapes, statistical face deformation models draw much attention, since they have numerous potential applications ranging from gaming to surveillance systems. However, the accuracy of statistical models is currently not sufficient to satisfy requirements such as the high geometric and visual accuracy needed for movie production. As a result, there is a definite need to obtain an accurate statistical model that can integrate all critical aspects from geometric and colour information to other appearance information at high resolution and with high accuracy.

Participants

- Mirela Ben-Chen
Technion – Haifa, IL
- Sébastien Bougleux
Caen University, FR
- Michael Breuß
BTU Cottbus, DE
- Alex M. Bronstein
Tel Aviv University, IL
- Michael M. Bronstein
University of Lugano, CH
- Alfred M. Bruckstein
Technion – Haifa, IL
- Leila De Floriani
University of Genova, IT
- Anastasia Dubrovina
Technion – Haifa, IL
- Maurizio Falcone
University of Rome "La Sapienza", IT
- Anna Hilsmann
HU Berlin, DE
- Yong Chul Ju
Universität Stuttgart, DE
- Margret Keuper
Universität Freiburg, DE
- Ron Kimmel
Technion – Haifa, IL
- Iasonas Kokkinos
Ecole Centrale – Paris, FR
- Aritom Kovnatsky
University of Lugano, CH
- Petros Maragos
National TU – Athens, GR
- Gerlind Plonka-Hoch
Universität Göttingen, DE
- Martin Rumpf
Universität Bonn, DE
- Frank R. Schmidt
Universität Freiburg, DE
- Carola-Bibiane Schönlieb
University of Cambridge, GB
- William Smith
University of York, GB
- Sibel Tari
Middle East Technical University – Ankara, TR
- Silvia Tozza
University of Rome "La Sapienza", IT
- Matthias Vestner
TU München, DE
- Martin Welk
UMIT – Hall in Tirol, AT
- Stefanie Wuhrer
Universität des Saarlandes, DE
- Hao Zhang
Simon Fraser University – Burnaby, CA
- Jovisa Zunic
University of Exeter, GB