Report from Dagstuhl Seminar 11142

Innovations for Shape Analysis: Models and Algorithms

Edited by

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

1 **Universität des Saarlandes, DE**
2 **Technion – Haifa, IL, freddy@cs.technion.ac.il**
3 **National TU – Athens, GR, maragos@cs.ntua.gr**

Abstract

This report documents the program and the results of Dagstuhl Seminar 11142 *Innovations for Shape Analysis: Models and Algorithms*, taking place April 3-8 in 2011. The focus of the seminar was to discuss modern and emerging topics in shape analysis by researchers from different scientific communities, as there is no conference specifically devoted to this field.

Seminar 03.–08. April, 2011 – www.dagstuhl.de/11142
1998 ACM Subject Classification I.4 Image Processing and Computer Vision
Keywords and phrases Shape analysis, mathematical morphology, shape reconstruction, numerical computing, level set methods, fast marching methods
Digital Object Identifier 10.4230/DagRep.1.4.23
Edited in cooperation with Silvano Galliani

1 Executive Summary

Michael Breuß
Alfred M. Bruckstein
Petros Maragos

The notion of *shape* is fundamental in image processing and computer vision, as the shape of objects allows the semantic interpretation of image contents. This is also known from human vision, as humans can recognise characteristic objects solely from their shapes. By *shape analysis* one denotes models and algorithms for detection and processing of shapes in images. It is at the heart of a lot of applications in sciences and engineering.

While the visual quality of an image benefits from a large number of pixels and a high resolution of details, the same assertion does not generally hold for important shape information. As a simple example, when refining an image of $N = n \times n$ pixels towards a higher resolution, the number of pixels in the contour of an object grows just linearly with n, but the image size N grows quadratically. This *linear scaling property* of shape descriptors is an attractive feature of shape analysis methods when considering large images.

Thanks to technological progress made within the last decade one can nowadays acquire high-resolution images with relatively inexpensive, common devices. An important example are digital cameras that allow to acquire images of several megapixels. The size of the image files has grown accordingly. As the trend of developing even more accurate and inexpensive acquisition devices will certainly continue in the next years, the linear scaling property of shape descriptors makes shape analysis methods an even more useful tool in image processing than in the past. However, there are also substantial *new challenges* in shape analysis: Concerning algorithms that allow the processing of the arising large data sets in acceptable time, and with respect to adequate shape analysis models that allow for an efficient algorithmic formulation.

The purpose of this seminar was to meet these challenges by bringing together researchers that are engaged in recent and upcoming developments in shape analysis models and numerical computing. As an example, the field of differential geometry has grown to be important for shape analysis during the last years, while a field like deformable shape modelling just begins to influence shape analysis methods. On the algorithmical side, there are many recent innovations that can be important for shape analysis. As examples, let us mention new broadly applicable, efficient fast marching schemes, or graph-based iterative algorithms. The individual areas in shape analysis and numerical computing share an interest in the described techniques. However, modelling is seen as a hot topic in computer science, while numerical computing is often seen as a mathematical domain. Also the various areas within shape analysis research can benefit from the discussion of models and methods that are modern in their respective fields.

The purpose of bringing together researchers from different disciplines was to explore the benefits of a *cross-disciplinary* point of view.

- Researchers in continuous-scale shape analysis brought their knowledge of differential and variational models and the related methods to the meeting.
- Researchers in discrete shape analysis brought to the meeting their knowledge about the latest techniques in graph-based shape analysis, discrete topology and related optimisation methods.
- Researchers in numerical computing brought to the meeting their knowledge of numerical techniques and of numerical analysis.

As the demands in the individual fields of shape analysis are high, the research grous in which the most interesting techniques are under development are quite specialised. Because of this, there is no regular conference or workshop that serves as a meeting place for an exchange of ideas of these groups.

The seminar was conducted in a conference style, where every contributor gave a talk of about 20 to 25 minutes. There was much time for extensive discussions in between the talks and in the evenings, and as documented by the very positive evaluation there was generally a very open and constructive atmosphere. While it is at the moment this report is written very difficult to identify a new fundamental aspect of shape analysis as a result of the workshop, lots of interesting aspects were discussed. As we believe, these will inspire novel developments in both modelling and algorithms.

2 Table of Contents

3 Overview of Talks

3.1 Geodesic Variations: Algorithms and Applications

Fethallah Benmansour (EPFL – Lausanne, CH)

Joint work of Benmansour, Fethallah; Santambrogio, Filippo; Carlier, Guillaume; Peyré, Gabriel

The computation of geodesic distance and minimal paths is an essential building block for many techniques and applications in computer vision, computer graphics and imaging. In all these fields, it makes sense to consider variational optimization problems that takes into account geodesic distances. This includes, for instance optimal sampling of surfaces, shape optimization, traffic congestion, travel time seismic imaging, surface reconstruction, etc. Several efficient numerical solvers are available for the computation of geodesic distances.

However, the numerical computation of the derivative of this distance with respect to some parameters has been much less investigated. In this talk I will first review the Sub-Gradient Marching Algorithm, which is a numerical procedure to evaluate the derivative of the geodesic distance with respect to the metric on regular grids. I will detail a novel class of algorithms that extend this work to triangulated mesh, and can compute derivative with respect to the metric and to the seeding points. I will show some applications to shape optimization, traffic congestion, and geodesic centroidal Voronoi tessellation.

3.2 Refined Homotopic Thinning Algorithms and Quality Measures for Skeletonisation Methods

Michael Breuß (Universität des Saarlandes, DE)

Topological skeletons are shape descriptors that have been applied successfully in practical applications. However, many skeletonisation methods lack accessibility, mainly due to the need for manual parameter adjustment and the shortage of tools for comparative analysis.

In our contribution we address these problems. We propose two new homotopy-preserving thinning algorithms: Flux-ordered adaptive thinning (FOAT) extends existing flux-based thinning methods by a robust automatic parameter adjustment, maximal disc thinning (MDT) combines maximal disc detection on Euclidean distance maps with homotopic thinning. Moreover, we propose distinct quality measures that allow to analyse the properties of skeletonisation algorithms. Tests of the new algorithms and quality assessment tools are conducted on the widely used shape database CE-Shape-1.

3.3 Local and Global Diffusion Geometry in Non-rigid Shape Analysis

Alex M. Bronstein (Tel Aviv University, IL)

Diffusion geometry, scale-space analysis, and study of heat propagation on manifolds have recently become a popular tool in data analysis in a variety of applications. In this talk, we will explore the applications of diffusion geometry to the problems of non-rigid shape representation, comparison, and retrieval. We will show that diffusion processes allow defining both local and global geometric structures. Local shape descriptors based on diffusion kernels allow representing shapes as collections of geometric "words" and "expressions" and approaching shape similarity as problems in text search and matching. Global structures are diffusion metrics, insensitive to shape deformations and topological changes. Representing shapes as metric spaces endowed with diffusion distances, we can pose the problem of shape similarity as a comparison of metric spaces using the Gromov-Hausdorff distance. As examples of applications we will show large-scale shape retrieval, correspondence computation, and detection of intrinsic symmetries in non-rigid shapes.

3.4 Intrinsic Symmetry

Michael M. Bronstein (Universität Lugano, CH)

Joint work of Bronstein, Michael M.; Raviv, Dan; Bronstein, Alexander; Kimmel, Ron; Hooda, Amit; Horaud, Radu

Symmetry and self-similarity is the cornerstone of Nature, exhibiting itself through the shapes of natural creations and ubiquitous laws of physics. Since many natural objects are symmetric, the absence of symmetry can often be an indication of some anomaly or abnormal behavior. Therefore, detection of asymmetries is important in numerous practical applications, including crystallography, medical imaging, and face recognition, to mention a few. Conversely, the assumption of underlying shape symmetry can facilitate solutions to many problems in shape reconstruction and analysis.

Traditionally, symmetries are described as extrinsic geometric properties of the shape. While being adequate for rigid shapes, such a description is inappropriate for non-rigid ones. Extrinsic symmetry can be broken as a result of shape deformations, while its intrinsic symmetry is preserved.

In this talk, we will present a generalization of symmetries for non-rigid shapes and two different numerical framework for their detection and classification.

3.5 Video analysis: a Tool Towards Unsupervised Learning of Shapes

Thomas Brox (Universität Freiburg, DE)

Unsupervised learning requires a grouping step that defines which data belong together. A natural way of grouping in images is the segmentation of objects or parts of objects. While pure bottom-up segmentation from static cues is well known to be ambiguous at the object level, the odds are much better as soon as objects move. I will present a method that uses long term point trajectories based on dense optical flow. Defining pair-wise distances between these trajectories allows to cluster them, which results in temporally consistent segmentations of moving objects in a video shot. In contrast to multi-body factorization, points and even whole objects may appear or disappear during the shot.

3.6 From MLS to Overparametrized Nonlocal Variational Methods

Alfred M. Bruckstein (Technion – Haifa, IL)

This talk discusses a variational methodology, which involves locally modeling of data from noisy samples, combined with semi-local or global model-parameters regularization. We show that this methodology yields many previously proposed algorithms, from the celebrated moving least squares (MLS) methods to the globally optimal over-parametrization methods recently published for smoothing and optic flow estimation. The unified look at the range of problems and methods previously considered also suggests a wealth of novel global functionals and local modeling possibilities. Moreover, the proposed non-local variational functional provided by this methodology greatly improves the robustness and accuracy compared to previous methods. Therefore the proposed methodology may be viewed as a basis for a general framework for a variety of problem domains in signal and image processing and analysis, such as denoising, adaptive smoothing, reconstruction and segmentation.

3.7 A Semi-Lagrangian Scheme for the AMSS Model

Elisabetta Carlini (University of Rome "Sapienza", IT)

Joint work of Carlini, Elisabetta; Ferretti, Roberto

I will present a Semi-Lagrangian scheme for the approximation of the Mean Curvature Motion power 1/3, better known in image community as Affine Morphological Scale Space (AMSS) model. I will analyse then the properties of the scheme: consistency, monotonicity and convergence. I will finally present some numerical test for image de-noising.

3.8 Modeling Morse Complexes in Arbitrary Dimensions

Leila De Floriani (University of Genova, IT)

Ascending and descending Morse complexes, determined by a scalar field f defined over a manifold M, induce a subdivision of M into regions associated with the critical points of f. In this talk, we present two simplification operators, called removal and contraction, for Morse complexes, which work in arbitrary dimensions. Together with their inverse refinement operators, we show that they form a minimally complete set of atomic operators to create and update Morse complexes on M. We also present a compact dimension-independent graph-based representation for Morse complexes which can be coupled with a discrete representation of the field as a simplicial mesh. We describe the effect of the simplification operators on such representation. Finally, we show some results and discuss current work and future perspectives.

3.9 Shape Analysis Problems in Practical Applications

Stephan Didas (Fraunhofer ITWM – Kaiserslautern, DE)

In this talk, we are sketching some shape analysis and characterisation problems motivated by current projects of the image processing department at Fraunhofer ITWM. The goals of these projects include automatic quality assurance in the production line or characterisation of different newly developed materials like fiber-reinforced polymers in material sciences, for example.In quality assurance applications, requirements on the system's response time often severely limit the choice of algorithms useable for the given task. Nevertheless, the presence of texture on the objects under investigation or the aim of allowing for a very general setting for the imaging process necessitate sophisticated methodologies to obtain the desired information from the image data. In material sciences, the measurement techniques are sometimes driven to their extremes (for example, in terms of spatial resolution). Thus the images can only have a very limited quality which makes a reliable analysis more complicated. Therefore, the intention of the presentation is not to present working solutions. We rather point out some (according to our knowledge) open problems from a practical point of view.

3.10 Non-rigid Shape Correspondence by Matching Pointwise Surface Descriptors and Metric Structures

Anastasia Dubrovina (Technion – Haifa, IL)

Finding a correspondence between two non-rigid shapes is one of the cornerstone problems in the field of three-dimensional shape processing. We describe a framework for marker-less non-rigid shape correspondence, based on matching intrinsic invariant surface descriptors, and the metric structures of the shapes. The matching task is formulated as a quadratic optimization problem that can be used with any type of descriptors and metric, and solved using an integer optimization tool. Further, we show the correspondence ambiguity problem arising when matching intrinsically symmetric shapes using only intrinsic surface properties. We show that when using isometry invariant surface descriptors based on eigendecomposition of the Laplace-Beltrami operator, it is possible to construct distinctive sets of surface descriptors for different possible correspondences. When used in a proper minimization problem, those descriptors allow us to explore number of possible correspondences between two given shapes.

Finally, we describe a hierarchical framework for an efficient solution of the quadratic matching problem.

3.11 Semi-Lagrangian Schemes for Nonlinear PDEs in Image Processing

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

The PDE approach to the solution of image processing problems has been rather successful in the last 30 years opening new directions in the field. Typically the PDEs appearing in some classical problems like Shape from Shading, segmentation and filtering are nonlinear and often degenerate so that standard numerical methods for their approximation can not be applied. We describe a class of numerical schemes for first and second order nonlinear PDEs based on semi-Lagrangian (SL) techniques. Those methods have been originally applied for simple advection equations, more recently they have been extended to deal with a wide range of nonlinear possibly degenerate differential problems. The main advantage is that SL schemes have strong stability properties and allow for large time steps still providing accurate solutions. We will give a short review of these properties as well as some applications to some classical image processing problems.

3.12 Geodesic Regression on Shape Manifolds

P. Thomas Fletcher (University of Utah, US)

In this talk I present a regression method, called geodesic regression, for modeling the relationship between a manifold-valued random variable and a real-valued independent parameter. The principle is to fit a geodesic curve, parameterized by the independent parameter, that best fits the data. Error in the model is evaluated as the sum-of-squared geodesic distances from the model to the data, and this provides an intrinsic least squares criterion. Geodesic regression is, in some sense, the simplest parametric model that one could choose, and it provides a direct generalization of linear regression to the manifold setting. I will also present a hypothesis test for determining the significance of the estimated trend. While the method can be generally applied to any form of manifold data, I will show a specific example of analyzing shape changes in the corpus callosum due to age.

3.13 Statistical Shape Models in Biomedical Image Segmentation and Visualization of Fuzzy Shapes

Hans-Christian Hege (ZIB – Berlin, DE)

In the first part of the talk a toolset is sketched for creation of 3D statistical shape models (including semi-automatic segmentation, geometry processing for creation of training shapes, flexible establishment of correspondences and PCA in Euclidean space) and segmentation of image data using various methods for adaptation of the model to new image data (including analysis of 1D intensity profiles along normal vectors or omni-directional, formulation of the optimization as MRF and computing the optimum using FastPD). Furthermore, the creation and utilization of articulated statistical shape models for modeling of anatomical joints is addressed.

In the second part of the talk the propagation of uncertainties (errors, noise) in a scalar field to positional uncertainty of level sets is discussed. This includes the role of numerical condition as well as computation of spatial probabilities that a given element (edge, face, ...) of a cell of the sampling grid crosses a level set of a given threshold. Two cases are discussed: uncorrelated random variables with arbitrary distribution and correlated Gaussian random variables. Examples from engineering and climate research illustrate the results.

3.14 Segmentation and Skeletonization on Arbitrary Graphs Using Multiscale Morphology and Active Contours

Petros Maragos (National TU – Athens, GR)

Joint work of Maragos, Petros; Kimon Drakopoulos

In this chapter we focus on formulating and implementing on abstract structures such as arbitrary graphs popular methods and techniques developed for image analysis, in particular multiscale morphology and active contours. To this goal we extend existing work on graph morphology to multiscale dilation and erosion and implement them recursively using level sets of functions defined on the graph's nodes. We propose approximations to the calculation of the gradient and the divergence of vector functions defined on graphs and use these approximations to apply the technique of geodesic active contours for object and edge detection. Finally, using these novel ideas, we propose a method for multiscale shape skeletonization on arbitrary graphs.

3.15 A PDE approach to Photometric Shape from Shading

Roberto Mecca (University of Rome "Sapienza", IT)

The Shape from Shading problem based on a single light source suffers from the convex/concave ambiguity. One possible solution to this limitation is to use two or more light sources, considering the Shape from Shading Photometric Stereo (SfS-PS) problem. In the talk I will introduce the SfS-PS problem and present some new theoretical results showing uniqueness. In fact, the use of two or more images based on different light sources allows to determine, under some hypotheses, the surface we are observing without additional information other than the light sources direction. I will also describe some numerical schemes based on the semi-lagrangian approximation of the new differential problem, discussing their efficiency and accuracy and comparing these methods with some classical finite difference schemes. Several tests on real and virtual images will be presented.

3.16 Discrete surface processing and adaptive remeshing

Serena Morigi (University of Bologna, IT)

We use various partial differential equation (PDE) models to efficiently solve several surface processing, including reconstruction, smoothing, and remeshing.

In particular, we propose a new adaptive remeshing strategy for the regularization of arbitrary topology triangulated surface meshes. Unlike existing sophisticated parametric remeshing techniques, our explicit method redistributes the vertices on the surface by keeping all edges on element stars approx-imatively of the same size, and areas proportional to the surface features. At this aim we solve a two-step PDE model using discrete differentialgeometry operators suitably weighted to preserve surface curvatures and to obtain a good mesh quality, that is well-shaped triangles. Several examples demonstrate that the pro-posed approach is simple, efficient and gives very desirable results especially for surface models having sharp creases and corners.

3.17 A truly Unsupervised, Non-Parametric Clustering Method

Pablo Muse (Universidad de la Republica – Montevideo, UY)

Joint work of Muse, Pablo; Tepper, Mariano; Almansa, Andrés

Human perception is extremely adapted to group similar visual objects. The Gestalt school studied the perceptual organization and identified a set of rules that govern human perception. One of the earlier and most powerful qualities, or gestalts, is proximity, which states that spatial or temporal proximity of elements may induce to perceive them as a single group. From an algorithmic point of view, the main problem with the gestalt rules is their qualitative nature. Our goal is to design a clustering method that can be considered a quantitative assessment of the proximity gestalt. We show that this can be achieved by analyzing the inter-point distances of the Minimum Spanning Tree, a structure that is closely related to human perception. We present a method that relies on the sole characterization of non-clustered data, thus being capable of detecting non-clustered data as such, and to detect clusters of arbitrary shape.

The method is fully unsupervised in the sense that the user input only relates to the nature of the problem to be treated, and not the clustering algorithm itself. Even the number of clusters does not need to be previously chosen.

Strictly speaking the method involves one single parameter that controls the degree of reliability of the detected clusters. However, the algorithm can be considered parameter-free, as the result is not sensitive to its value.

3.18 Image-based 3D Modeling via Cheeger Sets

Martin Oswald (TU München, DE)

Joint work of Oswald, Martin; Toeppe, Eno; Cremers, Daniel; Rother, Carsten

In this talk, we present a novel variational formulation for generating $3D$ models of objects from a single view. Based on a few user scribbles in an image, the algorithm automatically extracts the object silhouette and subsequently determines a $3D$ volume by minimizing the weighted surface area for a fixed user-specified volume. The respective energy can be efficiently minimized by means of convex relaxation techniques, leading to visually pleasing smooth surfaces within a matter of seconds. In contrast to existing techniques for single-view reconstruction, the proposed method is based on an implicit surface representation and a transparent optimality criterion, assuring high-quality $3D$ models of arbitrary topology with a minimum of user input.

3.19 Nested Sphere Statistics of Skeletal Models

Stephen M. Pizer (University of North Carolina – Chapel Hill, US)

We seek a form of object model that exactly and completely captures the interior of most non-branching anatomic models and simultaneously is well suited for probabilistic analysis on populations of such objects. We show that certain nonmedial, skeletal models satisfy these requirements. These models will first be mathematically defined in continuous 3-space, and then discrete representations will be derived. We will describe means of fitting skeletal models into manual or automatic segmentations of objects in a way stable enough to support statistical analysis, and we will specify means of modifying these fits to provide good correspondences of spoke vectors across a training population of objects. Understanding will be developed that these discrete skeletal models live in an abstract space made of a Cartesian product of a Euclidean space and a collection of spherical spaces. Based on this understanding and the way objects change under various rigid and nonrigid transformations, a method analogous to principal component analysis called composite principal nested spheres will be seen to apply to learning an efficient collection of modes of object variation about a Fréchet mean object.

The methods will be illustrated by application to a few anatomic objects.

3.20 Group-valued regularization for motion segmentation of dynamic non-rigid shapes

Guy Rosman (Technion – Haifa, IL)

Understanding of articulated shape motion plays an important role in many applications in the mechanical engineering, movie industry, graphics, and vision communities. In this paper, we study motion-based segmentation of articulated 3D shapes into rigid parts. We pose the problem as finding a group-valued map between the shapes describing the motion, forcing it to favor piecewise rigid motions. Our computation follows the spirit of the Ambrosio-Tortorelli scheme for Mumford-Shah segmentation, with a diffusion component suited for the group nature of the motion model. Experimental results demonstrate the effectiveness of the proposed method in non-rigid motion segmentation.

3.21 Variational Models in Shape Space and Links to Continuum Mechanics

Martin Rumpf (Universität Bonn, DE)

The analysis of shapes as elements in a frequently infinite-dimensional space of shapes has attracted increasing attention over the last decade. The aim of this talk is to adopt a primarily physical perspective on the space of shapes and to relate this to the prevailing geometric perspective. Indeed, we consider shapes given as boundary contours of volumetric objects, which consist either of an elastic solid or a viscous fluid. In the first case, shapes are transformed via elastic deformations, where the associated elastic energy only depends on the final state of the deformation and not on the path along which the deformation is generated. The minimal elastic energy required to deform an object into another one can be considered as a dissimilarity measure between the corresponding shapes. We apply this approach for shape averaging and shape statistics. In the second case, shapes are transformed into each other via viscous transport of fluid material, and the flow naturally generates a connecting path in the space of shapes. The viscous dissipation rate, the rate at which energy is converted into heat due to friction, can be defined as a metric on an associated Riemannian manifold.

3.22 3D Curve Skeleton Computation and Use for Discrete Shape Analysis

Gabriella Sanniti Di Baja (CNR – Naples, IT)

A discrete 3D curve skeletonization algorithm is described, based on iterated voxel removal guided by the distance transform of the object and on anchor point selection. The anchor point selection criterion is adequate to originate a curve skeleton reflecting object's shape sufficiently well. Then, the use of the curve skeleton for object decomposition in the framework of the structural approach to shape analysis is discussed. A suitable partition of the skeleton is presented that originates object decomposition in accordance with human intuition.

3.23 Incremental Level Set Tracking

Nir Sochen (Tel Aviv University, IL)

We consider the problem of contour tracking in the level set framework. Level set methods rely on low level image features, and very mild assumptions on the shape of the object to be tracked. To improve their robustness to noise and occlusion, one might consider adding shape priors that give additional weight to contours that are more likely than others. This works well in practice, but assumes that the class of object to be tracked is known in advance so that the proper shape prior is learned. In this work we propose to learn the shape priors on the fly. That is, during tracking we learn an eigenspace of the shape contour and use it to detect and handle occlusions and noise. Experiments on a number of sequences reveal the advantages of our method.

3.24 Global Minimization for Continuous Multiphase Partitioning Problems Using a Dual Approach and graph cuts algorithms

Xue-Cheng Tai (University of Bergen, NO)

Joint work of Tai, Xue-Cheng; Baue, Egil; Yuan, Jing

This talk is devoted to the optimization problem of continuous multi- partitioning, or multi-labeling, which is based on a convex relaxation of the continuous Potts model. In contrast to previous efforts, which are trying to tackle the optimal labeling problem in a direct manner, we first propose a novel dual model and then build up a corresponding duality-based approach. By analyzing the dual formulation, we can show that the relaxation is often exact, i.e. the optimal solution is also globally opti- mal to the nonconvex Potts model. In order to deal with the nonsmooth dual problem, we suggest a smoothing method based on the log-sum ex- ponential function and also indicate that such smoothing approach gives rise to the novel smoothed primal-dual model and suggests labelings with maximum entropy. Such smoothing method for the dual model produces an expectation maximization algorithm for the multi-labeling problem. Numerical experiments show competitive performance in terms of quality and efficiency compared to several state of the art methods for the Pott's model. In the end, we will also present several recent algorithms for computing global minimizers based on graph cut algorithms and augmented Lagrangian approaches.

3.25 Non-Local Ambrosio-Tortorelli and 3-Partite Skeletons

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

The phase field of Ambrosio-Tortorelli, which results from a competition between phase separation and local neighborhood interaction, has proven to be an indispensable tool in variational formulations of shape analysis that jointly involve region and boundary terms. When the local neighborhood interaction is emphasized over the phase separation, a smooth distance function, which codes local symmetries hence skeletons, is obtained.

In this talk, I will introduce a non-local modification to the otherwise local interaction term by adding the L^2 norm of the expectation, which suggests further regularization of the smooth distance function by a zero-sum function. After discussing the modified field, I will show alternative skeleton constructions leading to 3-partite forms separately coding bodies, appendages, and boundary texture.

3.26 Integrated DEM Construction and Calibration of Hyperspectral Imagery: A Remote Sensing Perspective

Christian Woehler (TU Dortmund, DE)

In this study we present an approach to the image-based construction of planetary digital elevation maps (DEM) by fusion of hyperspectral imagery with depth data obtained by laser altimetry measurements. Photometric methods like photoclinometry, shape from shading, or photometric stereo yield dense surface normal fields with a lateral resolution coming close to that of the images themselves, but the inferred relative depth data tend to be strongly biased on large scales, while laser altimetry provides absolute depth data with the drawback of high-frequency noise and limited lateral resolution. Our DEM construction algorithm applies an iterative scheme to combine absolute depth and gradient data. It initially operates on coarsely downscaled absolute depth data and then successively increases the resolution for each iteration step. During each step, the depth data and the image radiances on the respective spatial scale are used to estimate the non-uniform surface albedo, which then yields an updated DEM. The resulting DEM reveals fine surface details as its lateral resolution is comparable to that of the original images. We describe the surface reflectance by the physically motivated Hapke model, which is an analytical approximation to the radiative transfer equation.

Since for DEM construction we employ 85-band hyperspectral image data of 140 m lateral resolution provided by the Chandrayaan-1 Moon Mineralogy Mapper (M3) instrument, we are able to obtain for each hyperspectral pixel a reflectance spectrum normalised to a standard configuration of 30°incidence angle, 0°emission angle, and 30°phase angle, where in contrast to standard hyperspectral calibration approaches the detailed topography is taken into account. We point out the effect of the wavelength-dependent Hapke model based calibration on important features extracted from the normalised spectra (wavelength, relative depth, and width of the characteristic lunar iron absorption trough at 1000 nm; spectral ratio between 2817 and 2657 nm indicating water and/or hydroxyl ions).

However, a detailed inspection of the spectral features reveals pronounced residual dependences on topography-induced illumination variations, which are probably due to imperfections of the Hapke model of the order of one percent or less. Hence, we propose an empirical approach to learn these residual deviations from compositionally homogeneous surface regions displaying a large variety of slopes, such as the inner walls of small craters, based on the high-resolution DEM data. We show that our empirical correction strongly reduces topography-induced illumination effects on the reflectance spectra. As an ultimate refinement step, the corrected reflectances can be used to determine corrected albedo maps and DEM data.

Finally, we discuss open issues, specifically the physical modelling of topography-induced effects (e.g. by introducing a wavelength-dependent single-particle angular scattering function), and the fundamental problem of distinguishing between effects related to illumination vs. surface temperature at infrared wavelengths.

3.27 Stochastic Diffeomorphic Evolution and Tracking

Laurent Younes (Johns Hopkins University, US)

We will describe how shape evolution can be controlled by "diffeomorphic finite elements", or diffeons. In such an approach, diffeons, which are shape dependent vector fields, are combined linearly and integrated in a dynamical system to generate diffeomorphic motion. As such, they can be used in segmentation problems, where they lead (combined with gradient descent) to diffeomorphic active contours. They can also be generated as a stochastic process, leading to random shape evolution. We will discuss the definition of such random processes, and focus in particular on "Eulerian" definitions of the process (depending on the current shape location), in contrast with "Lagrangian" definitions that relate to fixed coordinate systems or parametrizations. We will show examples of shape generated by such processes and discuss preliminary applications in diffeomorphic shape tracking.

3.28 Shape Analysis for 3D Point Cloud.

Hong-Kai Zhao (University of California – Irvine, US)

In the first part I will present an efficient algorithm for computing the Euclidean skeleton of an object directly from a point cloud representation on an underlying grid. The key point of this algorithm is to identify those grid points that are (approximately) on the skeleton using the closest point information of a grid point and its neighbors. The three main ingredients of the algorithm are: (1) computing closest point information efficiently on a grid, (2) identifying possible skeletal points based on the number of closest points of a grid point and its neighbors with smaller distances, (3) applying a distance ordered homotopic thinning process to remove the non-skeletal points while preserving the end points or the edge points of the skeleton.

In the second part I will present our recent work on present implicit surface reconstruction algorithms for point clouds. We view the implicit surface reconstruction as a three dimensional binary image segmentation problem that segments the computational domain into an interior region and an exterior region while the boundary between these two regions fits the data points properly. The key points with using an image segmentation formulation are: (1) an edge indicator function that gives a sharp indicator of the surface location, and (2) an initial image function that provides a good initial guess of the interior and exterior regions. In this work we propose novel ways to build both functions directly from the point cloud data. We then adopt recent convexified image segmentation models and fast computational algorithms to achieve efficient and robust implicit surface reconstruction for point clouds.

3.29 Distance Images and Intermediate-Level Vision

Steven W. Zucker (Yale University, US)

The early stages of computer vision are dominated by image patches or features derived from them; high-level vision is dominated by shape representation and recognition. However there is almost no work between these two levels, which creates a problem when trying to recognize complex categories such as "airports" for which early feature clusters are ineffective. We argue that an intermediate-level representation is necessary and that it should incorporate certain high-level notions of distance and geometric arrangement into a form derivable from images.

We propose an algorithm based on a reaction-diffusion equation that meets these criteria; we prove that it reveals (global) aspects of the distance map locally; and illustrate its performance on airport and other imagery, including visual illusions. Finally, we observe that the feedforward and feedback pathways that define the intermediate-levels of biological vision systems could benefit directly from such models, and we sketch one plausible path for implementing them via local field potentials.

3.30 Orientation and Anisotropy of Multicomponent Shapes

Jovisa Zunic (University of Exeter, GB)

Orientation and anisotropy of multicomponent shapes will be considered. There are many situations where a number of single objects are better considered as components of a multicomponent shape (e.g. a fish shoal), but also there are situations where a single object is better segmented into natural components and considered as a multicomponent shape (decomposition of cellular materials onto the corresponding cells). These problems have not been considered previously, even though both orientation and anisotropy problems of single component shapes have been considered earlier.

Participants

- Fethallah Benmansour
EPFL - Lausanne, CH
- Michael Breuß
Univ. des Saarlandes, DE
- Alex M. Bronstein
Tel Aviv University, IL
- Michael M. Bronstein
Universität Lugano, CH
- Thomas Brox
Universität Freiburg, DE
- Alfred M. Bruckstein
Technion – Haifa, IL
- Elisabetta Carlini
University of Rome "Sapienza", IT
- Leila De Floriani
University of Genova, IT
- Stephan Didas
Fraunhofer ITWM – Kaiserslautern, DE
- Anastasia Dubrovina
Technion – Haifa, IL
- Maurizio Falcone
University of Rome "Sapienza" IT
- P. Thomas Fletcher
University of Utah, US
- Silvano Galliani
Univ. des Saarlandes, DE
- Hans-Christian Hege
ZIB – Berlin, DE
- Petros Maragos
National TU - Athens, GR
- Roberto Mecca
University of Rome "Sapienza" IT
- Serena Morigi
University of Bologna, IT
- Pablo Muse
Universidad de la Republica – Montevideo, UY
- Martin Oswald
TU München, DE
- Stephen M. Pizer
University of North Carolina – Chapel Hill, US
- Guy Rosman
Technion - Haifa, IL
- Martin Rumpf
Universität Bonn, DE
- Gabriella Sanniti Di Baja
CNR – Naples, IT
- Nir Sochen
Tel Aviv University, IL
- Xue-Cheng Tai
University of Bergen, NO
- Sibel Tari
Middle East Technical University – Ankara, TR
- Christian Wöhler
TU Dortmund, DE
- Laurent Younes
Johns Hopkins University, US
- Hong-Kai Zhao
University of California – Irvine, US
- Steven W. Zucker
Yale University, US
- Jovisa Zunic
University of Exeter, GB