

GMCI'2002 - Report
(published on *citr.auckland.ac.nz/dgt*)

11th Workshop “Theoretical Foundations of Computer Vision”

Geometry, Morphology and Computational Imaging

07-12 April 2002, Dagstuhl (Germany)

Image analysis and computer graphics depend on geometric modelling and analysis of objects in two- or multidimensional spaces. Different disciplines such as digital geometry, mathematical morphology, polyeder geometry or computational geometry, just to cite a few, are closely related to progress in image analysis and computer graphics.

The workshop discussed theoretical fundamentals related to those issues and specified open problems and major directions of further development in the field of geometric problems related to image analysis and computer graphics. The seminar schedule was characterized by flexibility, working groups, and sufficient time for focused discussions. There will be an edited volume of seminar papers (within the Springer LNCS series).

The contributions during the workshop have been related to one of the following subjects:

(1) geometric algorithms for image processing or computer vision for extracting structures from images, geometric shape matching, image segmentation and image restoration, or image halftoning,

(2) mathematical morphology (multiresolution representations, texture models, lattice-theoretical and fuzzy models),

(3) geometric feature analysis (length of a curve, area of surfaces in 3D, curvature), or

(4) further geometric aspects of computer vision or image processing occurring in image acquisition, optical illusions, shape recovery or depth analysis, or modelling of complex situations in vision-based robotics.

The 39 abstracts of this workshop report are in the same order as talks had been scheduled in the workshop program. The workshop had 41 participants: 10 from Japan, 9 from France, 5 from Germany, 3 from Israel, New Zealand and USA each, 2 from The Netherlands and Slovakia each, and one from Australia, Belgium, Canada and Italy each.

The participants enjoyed the atmosphere of the Dagstuhl center for Computer Science, which allowed a very productive workshop.

Tetsuo Asano, Reinhard Klette, and Christian Ronse
15 April 2002

Spatial Relationships between Objects and Fuzzy Objects using Mathematical Morphology

Isabelle Bloch (ENST CNRS URA 820, Paris)

Basic mathematical morphology operations rely mainly on local information, based on the concept of a structuring element. But mathematical morphology also deals with more global and structural information since several spatial relationships can be expressed in terms of morphological operations (mainly dilations). We consider topological (which include part-whole relationships such as inclusion, exclusion, adjacency, etc.) and metric relationships (distances and directional relative position). The interest in these relations is being highlighted in different areas such as vision, GIS, cognitive psychology, or artificial intelligence.

The aim of this presentation is to show that this framework allows to represent spatial relationships in a unified way in various settings: a purely quantitative one if objects are precisely defined, a semi-quantitative one if objects are imprecise and represented as spatial fuzzy sets, and a qualitative one, for reasoning in a logical framework about space. This is made possible thanks to the strong algebraic structure of mathematical morphology, that finds equivalents in set theoretical terms, fuzzy operations and logical expressions. The proposed framework allows us to address three questions:

(1) We consider the problem of defining and computing spatial relationships between two objects, in both crisp and fuzzy cases.

(2) Then we propose a way to represent spatial knowledge in the spatial domain. Given a reference object, we define a spatial fuzzy set that represents the area of the space where some relationship to this reference object is satisfied (to some degree). The advantage of these representations is that they map all types of spatial knowledge in the same space, which allows for their fusion and for spatial reasoning.

(3) Finally we show that spatial relationships can be expressed in the framework of normal modal logics, using morphological operations applied on logical formulas. This can be useful for symbolic (purely qualitative) spatial reasoning.

Morphological Texture Analysis using Optimization of Structuring Elements

Akira Asano (Hiroshima University)

This presentation proposes two methods of texture primitive description based on optimization of the shape of a structuring element that fits the shape of a primitive. We assume a model that a texture is composed of grains derived from one primitive with some fluctuation on shapes and sizes.

The first method [1] is based on an assumption on the size distribution of grains.

We assume that the target texture contains grains of various sizes so that the size density function of the texture is uniform. Thus the optimization of the shape of the structuring element for the measurement of size density function to make the size density function uniform yields the best estimate of the primitive. The simulated annealing method is employed for the optimization. The second method [2] requires no assumption on the size distribution. This method is based on the morphological skeleton. The skeleton of an object has a property that it is reduced to one point when the structuring element used for the skeletonization is exactly homothetic to the object. If we assume that the target texture is composed of grains derived from one primitive by the homothetic expansion, the structuring element that produces the skeleton of the minimum number of pixels is homothetic to the primitive of the texture. The simulated annealing is again employed for the minimization.

References:

- [1] A. Asano, M. Miyagawa, and M. Fujio: Texture modelling by optimal gray scale structuring elements using morphological pattern spectrum. Proc. 15th Internat. Conf. on Pattern Recognition, 3, 479–482 (2000).
- [2] A. Asano, T. Ohkubo, M. Muneyasu, and T. Hinamoto: Texture primitive description using morphological skeleton. Proc. Internat. Symp. on Mathematical Morphology VI, 101–108 (2002).
- [3] A. Asano, J. Endo, and C. Muraki: Multiprimitive texture analysis using cluster analysis and size density function. Proc. Internat. Symp. on Mathematical Morphology VI, 109–116 (2002).

**Root Images of Median Filters
- Semi-Topological Approach -**

Ulrich Eckhardt (Universität Hamburg)

Median filters are frequently used in signal analysis because of their smoothing properties and their insensitivity with respect to outliers in the data. Since median filters are nonlinear filters, the tools of linear theory are not applicable to them. One approach to deal with nonlinear filters consists in investigating their root images (fixed elements or signals transparent to the filter). Whereas for one-dimensional median filters the set of all root signals can be completely characterized, this is not true for higher dimensional filters.

In 1981 Tyan proposed a method for constructing “small” root images for two-dimensional median filters. In 1989, Döhler attempted to prove validity of Tyan’s construction. Although Döhler’s results are true for a wide class of median filters, his arguments were not correct and his assertions do not hold universally.

By imbedding the problem into the framework of semi-topology (Latecki, 1992)

and by applying convexity arguments it is possible to prove the correctness of the Tyan-Döhler construction for a wide class of median filters, it is also possible to extend it to the general d -dimensional case.

Multiresolution MIP Volume Rendering by Morphological Pyramids

Jos Roerdink (University of Groningen)

A multiresolution representation for maximum intensity projection (MIP) volume rendering is presented, based on morphological pyramids which allow progressive refinement and have the property of perfect reconstruction. In this approach the pyramidal analysis and synthesis operators are composed of morphological operators, combined with dyadic downsampling for analysis and dyadic upsampling for synthesis. Two types of pyramid are used, one in which the morphological operator during analysis is erosion, and another in which a morphological opening is used. For the second type of pyramid the approximation accuracy in higher levels of the pyramid is improved. Further improvements may be obtained by using still other types of morphological operators, based on the theory of connected operators.

Combinatorial and Geometric Problems Related to Digital Halftoning

**Tetsuo Asano (JAIST), Naoki Katoh (Kyoto University)
Koji Obokata (JAIST), and Takeshi Tokuyama (Tohoku University)**

The quality of color printers has been drastically improved in recent years, mainly based on the development of fine control mechanisms. On the other hand, there seems to be no great invention on the software side of the printing technology. What is required is a technique to convert a continuous-tone image into a binary image consisting of black and white dots that looks similar to the input image. Theoretically speaking, the problem is how to approximate an input continuous-tone image by a binary-tone image. Since this is one of the central techniques in computer vision and computer graphics, a great number of algorithms have been proposed with several theoretical results by the authors. The purposes of this are to reveal that there are a number of problems related to combinatorial and computational geometry and to present some solutions or clues to those problems.

Matrix Rounding under the L_p -Discrepancy Measure and Digital Halftoning

Naoki Katoh (Kyoto University), Tetsuo Asano (JAIST)
Koji Obokata (JAIST) and Takeshi Tokuyama (Tohoku Univ.)

In this paper, we study the problem of rounding a real-valued matrix into an integer-valued matrix to minimize an L_p -discrepancy measure between them. To define the L_p -discrepancy measure, we introduce a family F of regions (rigid submatrices) of the matrix, and consider a hypergraph defined by the family. The difficulty of the problem depends on the choice of the region family F . We first investigate the rounding problem by using integer programming problems with convex piecewise-linear objective functions, and give some nontrivial upper bounds for the L_p -discrepancy. In particular, when F is so called *unimodular*, the L_p -discrepancy is bounded by $|F|/2$ for $p \leq 3$. For $F =$ all 2×2 regions which are not unimodular, we prove the nontrivial upper bound $2|F|/3$. Then, we propose a “laminar family” for constructing a practical and well-solvable class of F . Indeed, we show that the problem is solvable in polynomial time based on the minimum-cost flow algorithm if F is a union of two laminar families. Finally, we show experimental results which suggest that the matrix rounding using L_1 -discrepancy for a union of two laminar families is suitable for developing a high-quality digital-halftoning software.

Reconstruction of Discrete Convex Sets from Point-source X-rays

Alain Daurat (University of Strasbourg)

A point-source X-ray gives the number of points on each semi-line originated from a fixed point. In this talk we show that seven aligned point-source X-rays determine discrete convex sets which do not intersect the alignment-line. For six aligned point-sources we can find a counter-example. On another hand we can always reconstruct a set which verifies a lower notion of convexity from point-source X-rays in polynomial time. All these results are in fact a simple generalization of the already known parallel-X-rays results. But the following conjecture is specific to point-source X-rays: a discrete convex set would be determined by three non-aligned sources. (The continuous version is a theorem of Volcic (1986)).

**Digital Sets and Their Polygonal Representation
by Corresponding Segments**

Helene Reiter (Universität Hamburg)

We discuss representations of 2D objects or images. The silhouette of an object can be depicted by its boundary curve. Objects can be recognized by way of the comparing shapes with given templates in a database. For efficiency it is important to simplify the silhouette of the object. We are able to approach the boundary curve of an object either by continuous or by discrete evolution. By way of the discrete evolution one has many advantages: there is no type change from the original image to the approximated image, the evolved image is trivially localized with respect to the original image, and the properties of the original image are identical with the evolved image. Significant properties in this sense are convexity and concavity of boundary pieces. In this work we propose an algorithm for simplification of an 8-boundary curve of a digital object. The algorithm specifies a polygonal representation with corresponding convexity and concavity properties. We make a list of all cases, that should be treated with it particularly. We use the Radon-transformation for the specification of the algorithm. The Radon-transformation (in the literature also known as Hough-transformation) is of basic importance for applications and fundamentals of image processing.

**Seven Open Problems
in Digital Geometry and Topology**

Reinhard Klette (The University of Auckland)

The talk discusses seven open questions in digital geometry and topology (see web site citr.auckland.ac.nz/dgt):

(1) Is there a method with superlinear convergence speed for estimating the length of digitized 2D curves?

(2) (Th. Bülow and R. Klette, 2001) Is there a linear-time algorithm for multigrid convergent estimation of the length of digitized 3D curves based on minimum-length polygonal curves?

(3) Provide a method for surface area estimation which is multigrid convergent (convergence theorem), calculates simple polyhedra for simple isothetic polyhedra, and possesses an efficient algorithm.

(4) (A. Rosenfeld and R. Klette, 2001) Does there exist a linear-time on-line algorithm for segmenting a digital curve into a minimal number of DSSs? Does there exist an efficient algorithm for segmenting a digital surface into a minimal number of DPSs? (as a start: Does it exist for surfaces of digitized convex sets?)

(5) (A. Imiya and U. Eckhardt, 2001) Improve the upper bound $2^n - 2^{n-2}$ for the number of possible digital topologies in Z^n , or show that this is the exact

number of digital topologies in Z^n .

(6) (Y. Kong, 2001) Construct a simplest possible theory that gives an axiomatic definition of ‘well-behaved 3D digital spaces (good pairs) and allows many results of 3D digital topology to be proved simultaneously for all such spaces.

(7) (Y. Kong and R. Klette, 2002) Consider simple polyhedra in 3D space. Is the problem of determining if two such polyhedra have isomorphic fundamental groups decidable or undecidable? If it is decidable then provide an algorithm to do so.

Digital Surface Area Estimation

David Coeurjolly, F. Flin and O. Teytaud
(Université Lumière Lyon 2)

Once fallen on the ground, snow undergoes a structure metamorphism governed by local temperature and humidity fields. Local 3D curvature is a governing parameter of metamorphism, whereas surface area denotes the ability for a given snow to evolve. To access these parameters from raw voxel data, the normal vector field has to be known all along the digitized surface of the snow sample. Meshing the initial B/W discrete image is the most popular approach; generally, it uses Marching Cubes methods: the easy derivation of normals from facets is used in many shading algorithms.

Many imaging scientists have developed normal vector estimation algorithms: a first approach consists in working on the volume of the image. Other algorithms use the surfel (surface element) decomposition of the digitized object, then average surfel elementary normals using convenient weightings and neighborhoods. An other interesting approach consists in considering two dimension slices of a 3D image.

As a precise vector field is computed on each surface voxel of the object, an accurate surface area estimation can be obtained directly from this field: when the surfel decomposition of the surface is available, this can be done by summing the projected area of each surfel along the local normal vector. A method applicable to raw voxel data is detailed. In order to analyze the behavior of the surface area estimator when the resolution of the grid increases, a proof of its multigrid convergence is also provided.

Digital Planes Identification and Recognition

Isabelle Sivignon (University of Grenoble)

With the evolution and improvement of image acquisition methods, we now have not only 2D images to study by also 3D volumes, that are for instance

often used in medical imaging. These 3D volumes are sets of voxels which inherent structure disables a good visualization as well as an easy storage, among other inconveniences. One solution is to transform these discrete objects in an invertible way into euclidean polyhedra.

Two main types of methods are usually used to solve this: using euclidean methods (for instance the marching cubes method) or using discrete methods. As we deal with discrete objects, discrete methods seem to be more appropriate for this problem. Many works (Papier and Franon in 1994, Debled in 1995, Papier in 1999, or Burguet and Malgouyres in 2000) have already been reported in this area, but in each case either the proposed algorithm finds a very huge number of faces, or it assumes some symmetric knowledge on the object, or often it is not invertible.

We propose here a polynomial algorithm that divides a 3D discrete object surface into patches of naive planes. This algorithm uses the naive plane recognition algorithm presented by Vittone in her PhD, which uses the dual space in order to get all the possible solutions for each recognized piece of plane. Using 8-neighborhood tracking of the boundary surfels, and allowing that planar patches may overlap, we obtain a decomposition that fits on simple examples (cube, chamfer cube, pyramid) what we visually would like to see, and where faces have regular shapes as any voxel of a piece of plane has at least three neighbors in this plane. We then point out in the algorithm the different choices that have to be done during voxel tracking and propose a new study on different strategies. The next step is now to find a coherent definition for discrete edges and vertices and then to define their position on the surface. But there we have to deal with some particular properties of naive plane intersections and another model may be required to achieve an invertible polyhedrization.

On Approximation and Representation of Surfaces in Implicit Forms

Fridrich Sloboda and Bedrich Zlatko (Slovak Academy of Sciences)

No abstract given.

Distance Map Based Image Enhancement for Interpolation

PeiFeng Zeng and Tomio Hirata (Nagoya University)

We propose an approach of distance map based image enhancement (DMIE) for improving the quality of interpolated images. In DMIE, edge detection is performed after images are interpolated by conventional interpolation schemes. A unified linear-time algorithm for distance transform is applied to deal with the calculation of Euclidean distance from pixels to edges. By warp calcula-

tion, pixels that are located around edges are translated towards edges. DMIE implements the sharpening of interpolated images with visual pleasant.

**Multiresolution Model
Based on Discrete Gradient Vector Fields**

**Leila De Floriani, Mohammed Mostefa Mesmoudi and
Emanuele Danovaro (University of Genova)**

In this work we introduce a discrete form of gradient vector field associated to a scalar field defined on vertices of a simplicial complex. We build up a decomposition of the complex similar to Smale decomposition corresponding to a Morse function defined on a manifold domain. In the case where the simplicial complex is a surface, we discuss the properties of critical points that allow us to classify them in four classes that generalize the nature of classical critical points of a Morse function. We define through our discussion the notion of a discrete Morse function. We apply our study to extract the main features of the domain and build up a multiresolution model representing the scalar field of the surface at different level of accuracy.

**Eyes from Eyes:
Analysis of Camera Design Using Plenoptic Video Geometry**

**Jan Neumann, Cornelia Fermüller, and Yiannis Aloimonos
(University of Maryland)**

We investigate the relationship between camera design and the problem of recovering the motion and structure of a scene from video data. The visual information that could possibly be obtained is described by the plenoptic function. A camera can be viewed as a device that captures a subset of this function, that is, it measures some of the light rays in some part of the space. The information contained in the subset determines how difficult it is to solve subsequent interpretation processes. By examining the differential structure of the time varying plenoptic function we relate different known and new camera models to the spatio-temporal structure of the observed scene. This allows us to define a hierarchy of camera designs, where the order is determined by the stability and complexity of the computations necessary to estimate structure and motion. At the low end of this hierarchy is the standard planar pinhole camera for which the structure from motion problem is non-linear and ill-posed. At the high end is a camera, which we call the full field of view polydioptric camera, for which the problem is linear and stable. In between are multiple-view cameras with large fields of view which we have built, as well as catadioptric panoramic sensors and other omni-directional cameras. We develop design suggestions for a

polydioptric camera, and based upon this new design we propose a linear algorithm for ego-motion estimation, which in essence combines differential motion estimation with differential stereo.

Morphology on Label Images

Christian Ronse (University of Strasbourg)

We consider morphological operations on images whose grey-levels or colors are considered as labels without ordering between them, except for a least element (meaning "no label") and a greatest element (meaning "any label" or "conflicting labels").

We investigate first flat morphological operators. This study is done in the framework of an arbitrary power lattice. Among the results in usual flat morphology for grey-level images, some extend in the general case, while some others don't, and for them we find counterexamples precisely in the lattice of label images.

Since the lattice of label images is not distributive, geodesical dilation and reconstruction can be properly defined only if we make some restrictions on the labels present in the mask image or in the marker image.

Given any connectivity on sets, there is a connection on label images, for which the connected components of an image are precisely its labeled flat zones. Furthermore, when the connectivity on sets arises from a symmetrical neighborhood function, the connected components of a label image can be reconstructed from markers through a geodesical reconstruction by dilation.

This variant of mathematical morphology can be used for the segmentation of moving objects in video sequences.

Quantification of the Convergence of Tissues in Lung for CAD of Chest X-ray CT Images

Yasushi Hirano (Nagoya University)

We propose two features to quantify the convergence of tissues (vessels, bronchus, interstitium, etc.) in lung for the computer aided diagnosis (CAD) using 3D chest X-ray CT images. The convergence is the phenomenon that tumor pulls the tissues surrounding the tumor toward the tumor. It is observed frequently in the malignant cases. Therefore the existence of convergence provides useful information for the benign/malignant discrimination. The convergence of vessels and interstitium are quantified in this study. The former is quantified by the concentration index, and the latter is quantified by the ratio of the volumes of Voronoi regions. The concentration index quantifies the concentration of line figures in 3-D space. If the linearized vessel regions are used as line figures, the

convergence of vessels is quantified. The concentration index takes the values between 0.0 and 1.0. If the line figures concentrate, the index takes a high value. The ratio of volumes of the Voronoi regions quantifies the difference of density of connected regions. Tumor pulls the interstitium near the tumor toward the tumor, if the convergence exists. This phenomenon causes the unbalance of the density of vessels. Comparing the volumes of Voronoi regions near the tumor and apart from tumor, the convergence of interstitium is quantified. The distance transformation and the thinning method are performed to the vessel region to get line figures that have radii of vessel region at each voxel. Voxels on the line figure are classified into groups according to their radii. A 3D Voronoi division is performed using a group. The feature is defined as the average volume of Voronoi regions in the lung over that near the tumor. The feature becomes about 1.0 for the case without convergence, and it becomes more than 1.0 for the case with convergence.

The experiments were performed using 78 cases of chest CT images. For 29 cases out of these 78 cases the degree of the convergence has been evaluated by medical doctors. Using these 29 cases, significance tests have been performed. We assumed that "There is no difference between the cases with convergence and cases without convergence" as the null hypothesis. The former feature and the latter feature are rejected with the level of significance 0.01 and 0.05, respectively. It is shown that the features are useful to quantify the degree of convergence. The result of the benign/malignant discrimination, the correct answer ratio is over 80%. It is confirmed that these features can be used for the benign/malignant discrimination.

k-Means Clustering and Applications in Image Processing

David M. Mount (University of Maryland)

Joint with T. Kanungo, N. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu

In k-means clustering we are given a set of n data points in d -dimensional space and an integer k , and the problem is to determine a set of k points, called centers, to minimize the mean squared distance from each data point to its nearest center. Clustering is a fundamental problem with many applications in image processing and quantization. No exact polynomial-time algorithms are known for this problem. Although asymptotically efficient approximation algorithms exist, these algorithms are not practical due to the extremely high constant factors involved. There are many heuristics that are used in practice, but we know of no bounds on their performance.

We consider the question of whether there exists a simple and practical approximation algorithm for k-means clustering. We present a local-search heuristic based on swapping centers in and out. We prove that this yields a $(9+\epsilon)$ -

approximation algorithm. We show that the approximation factor is almost tight, by giving an example for which the algorithm achieves an approximation factor of $(9-\epsilon)$. To establish the practical value of the heuristic, we present an empirical study that shows that, when combined with Lloyd's algorithm, this heuristic performs very well in practice on both synthetically generated data sets and data sets arising from color- and vector-quantization applications.

Representation of Plant Leaves using Surface Fitting Techniques

Birgit Loch (CPAI, University of Queensland)

Joint research with John Belward and Jim Hanan (CPAI)

Leaves play a vital role in plant development, serving as major resource collectors.

In this talk, we concentrate on modelling leaf surfaces for use in virtual plants, computer simulations that can serve as a valuable tool for biologists studying real plants. We first describe data acquisition using three different methods, i.e. sonic digitizer, laser scanner and photogrammetry. Then, we suggest interpolation methods for fitting leaf surfaces using piecewise polynomials on triangular grids. We explain how we deal with non-convex domains, adjusting the triangulation accordingly, and then apply an adaptive method that gives some control over the overall error when only a small number of data points is used. Finally, we display the resulting surface using texture mapping.

Bayesian Fitting

Michael Werman (University of Jerusalem)

It is common practice to fit parametric models (lines, circles, implicit polynomials etc) to data points, by minimizing the sum of squared distances from the points to the model (the MSE, or Mean Square Error, approach). While the MSE algorithm may seem natural, it in fact implicitly assumes that each data point is the noised version of the point on the model which is closest to it. This assumption is clearly false, and results in strong bias, for instance when fitting circles to data contaminated with strong noise.

The MSE algorithm suffers from another drawback: it cannot differentiate between a "large" model and a "small" one. For instance, when fitting a line segment to bounded data, one would like to know not only the slope and location of the fitted line, but also its start and end points. The MSE criterion cannot differentiate between the "correct" segment and a segment which is too long, because both have the same MSE error with respect to the data.

We offer a simple paradigm for fitting parametric models, which solves these

two problems. This is done by considering each point on the parametric model as a possible source for each data point. The model is also extended to non-parametric models, and gives good results even for curves with strong discontinuities.

Contour Shape Analysis Using Crystalline Flow

**Koichiro Deguchi (Tohoku University), and
Hidekata Hontani (Yamagata University)**

The crystalline flow is a simulation of crystal growth. It simulates anisotropic growths of crystals with non-local velocity. We apply it to realize the generation of a 2D crystal having a given contour shape. First, starting from the given contour shape, we shrink it by reversing the crystalline flow. In this reverse growth, we extract a time sequence of control parameters to lead the contour to a point seed. Then, using this sequence of parameters, we generate just the given contour shape. If we succeed in the regeneration of the shape, the parameters are considered to represent characteristics of the given contour shape.

The crystalline flow treats simple polygonal shape crystals. To represent the anisotropy of its growth, a so-called the *Wulff shape*, which is a convex polygon, is defined. All outward normals of facets of an applicable contour should belong to a set of normals of the Wulff shape, and the normal of adjacent facets to every facet should be parallel to the normal adjacent of the corresponding facet in the Wulff Shape. The crystal contour satisfying these conditions is said to be admissible, and proved to converge to a point seed by applying the reverse growth.

The growing velocity of every contour facet at every time is give by a function of its outward normal vector and generalized weighted curvature of the facet. This curvature of a facet is given as the reciprocal of its relative length to its corresponding facet length in the Wulff shape and a transition number. The transition number is +1 for a convex facet, -1 for concave, and 0 for otherwise. In the reverse process of the crystalline flow, first, concave facets and step shape disappear, and then all facets disappear into a point. So, the reverse growth is well described by a time sequence of facet disappearances and changes of the transition numbers in the process.

We propose to register this sequence on a plane spanned by the position on the original contour and the time of the reverse growth. Also, we propose a hierarchical representation to characterize a hierarchical relations between contour shapes. We illustrate generations of given contour shapes using registered data based on performed experiments.

Geometric Structure and Randomness in Texture Analysis and Simulation

Georgy L. Gimel'farb (The University of Auckland)

Different scenarios of probabilistic texture modelling based on Markov/Gibbs random fields with multiple pairwise pixel interactions are discussed.

Markov/Gibbs models consider textures as samples of a Markov random field with a joint Gibbs probability distribution of signals. One of the basic goals of modelling is to produce realistic large-size images after adapting the model to a particular texture using a small-size training sample. The model describes homogeneity of a texture in terms of repetitive geometric structure of local pixel interactions. Gibbs potentials specifying probabilities of signal co-occurrences in the interacting pixels govern random variations of the signals in different simulated images. For some regular mosaics, the model-based interaction map (MBIM) representing partial interaction energies for a large set of pixel pairs allows for selecting an orientation and size of rectangular repetitive patches that form the texture. These patches act as texels (texture elements), and their explicit specification considerably accelerates simulation of large-size prototypes of the mosaics.

Experiments with various natural mosaics show that the desired partitioning of the training sample into a set of rectangular texels can be obtained by analyzing positions of peripheral energy clusters in the training MBIMs, in particular, the positions of the clusters with the highest and second highest energy that do not occlude each other from the origin of the MBIM.

Voxel-based Surface Area Estimation: From Theory to Practice

Guy Windreich and Nahum Kiryati (Tel Aviv University)
G. Lohmann (Max Planck Inst. for Cognitive Neuroscience, Leipzig)

Consider a complex, highly convoluted three dimensional object that has been digitized and is available as a set of voxels. How can one estimate the (original, continuous) area of a region of interest on the surface of the object?

The problem appears in the analysis of segmented MRI brain data and in other three dimensional imaging applications. Several difficulties arise. First, due to the complexity of the surface and its foldings, the region of interest and its intended boundary can be concealed and are therefore difficult to delineate. Second, the correct surface topology on intricate voxel sets may not be obvious. Third, the surface area of a digital voxels world is generally very different than the area of the underlying continuous surface.

These problems can be partly circumvented by transforming the voxel data into a polyhedral surface representation. Our challenge is to accomplish the task while maintaining the original voxel representation. Estimators for the continu-

ous surface area of digital objects have been available for some time. However, the known methods are limited to fairly smooth and “well behaved” surfaces. This research bridges the gap between the available surface area estimation theory, that applies to idealized settings, and the reality of MRI brain data. Surface connectivity ambiguities are alleviated by considering the object/background boundary voxels rather than the border voxels themselves. The region of interest is delimited by growing geodesic paths between user-provided anchor points. Surface area estimation is extended to admit surfaces with higher curvature than previously considered. Performance evaluation results are provided, and operation on MRI brain data is demonstrated.

3D Border Tracking with Surface Structures

**Yukiko Kenmochi (ESIEE, Paris / JAIST) and
Atsushi Imiya (Chiba University)**

There already exist various border tracking algorithms of objects in a 2D/3D digital image. One of the basic and useful ideas to solve the border tracking problem in 2D is to use a curve structure of the border of a 2D object. A simple extension of the border tracking in 3D will be then to use a surface structure of the border of a 3D object. However, it is actually difficult to find a good definition or representation of 2D surfaces in a 3D discrete space such as a 3D digital image.

In the talk, we first give an overview and classification of various approaches for defining 2D discrete surfaces in a 3D discrete space. For our discrete surface representation which provide a good solution for the border tracking problem, we apply some basic ideas of combinatorial topology and algebraic topology, which are also used in previous work (but in different ways). After obtaining a discrete surface representation, we consider an algorithm for tracking the border which is always guaranteed to have a surface structure (but allowing non-manifold cases). We also give some interesting properties of our tracked border and make comparisons with results on other border tracking algorithms such as a simple boundary operation in mathematical morphology.

Junction Detection and the Topological Structure of Images

Ullrich Köthe (Universität Hamburg)

Segmentation is one of the fundamental tasks in image analysis. In order to arrive at a satisfying segmentation result, both bottom-up and top-down algorithms complementing each other will be needed. Unfortunately, bottom-up algorithms do not work very well in practice. This leads to a gap between low-level and high-level processing that is hard to bridge with high-level expect-

tations alone. Consequently, each segmentation method is restricted to a narrow application domain, and failures are quite common.

In this contribution we investigate some of the reasons for the failures of bottom-up processing. Three critical issues are identified: first, the segmentation problem has not been sufficiently split into separate subproblems which can be treated one at a time. Second, there is a lack in ground truth so that different algorithms cannot be objectively compared, and different types of failures cannot be attributed to different causes in a definite way. Third, there is no realistic theory about fundamental limits of bottom-up processing, and it is not known whether there exist algorithms that are optimal in the sense that they achieve the theoretical bounds.

In order to overcome these problems, we propose to use piecewise constant scenes as a first benchmark. These scenes are viewed through a real camera or subjected to a realistic camera model (including the effects of a point spread function, spatial sampling, noise and quantization). Although this scenario is a strong simplification of reality, it is still sufficiently rich in the sense that it contains many potential sources of error. Since topology is a more fundamental property of a scene than geometry, failures are measured in terms of how well the segmentation recovers the topological structure of the original scene. A hypothesis is discussed that algorithms which fail on these simplified images will likely fail on more realistic problems as well.

A test bed is described that realizes the proposed analysis of piecewise constant scenes. Initial investigations of popular segmentation algorithms (including zero crossings of the Laplacian of Gaussian, Canny's algorithm, and the watershed algorithm) reveal that none of these algorithms performs satisfactorily near corners and junctions. A number of topology violations, such as vertex splitting, spurious or broken edges, and phantom regions, regularly occur even without any noise. These results are in line with previous findings of similar studies. The analysis suggests several ways to overcome the problems.

Uncertainty in Visual Processes Predicts Geometrical Optical Illusions

**Cornelia Fermüller and Yiannis Aloimonos
(University of Maryland)**

It is proposed that many geometrical optical illusions, as well as illusory patterns due to motion signals in line drawings, are due to the statistics of visual computations. The interpretation of image patterns is preceded by a step where image features such as lines, intersections of lines, or local image movement must be derived. However, there are many sources of noise or uncertainty in the formation and processing of images, and they cause problems in the estimation of these features; in particular, they cause bias. As a result, the locations of

features are perceived erroneously and the appearance of the patterns is altered. The bias occurs with any visual processing of line features; under average conditions it is not large enough to be noticeable, but illusory patterns are such that the bias is highly pronounced. Thus, the broader message of this work is that there is a general uncertainty principle which governs the workings of vision systems, and optical illusions are an artifact of this principle.

Panoramic Imaging - Theory and Practice

Fay Huang, Shou-kang Wei, and Reinhard Klette
(The University of Auckland)

The general camera model of a cylindrical panoramic image is introduced. The model consists of four parameters. In particular, the radius of rotation R and the viewing angle ω of a line camera are the two main parameters for characterizing a panoramic view. Our camera model also enables us to discuss the generation of pairs of stereo panoramas.

We show how the 3D space is sampled by a pair of stereo panoramas by analyzing the following topics: sampling resolution, sampling distribution, and sampling distances. A major conclusion of our study on spatial sampling is that the parameter R has no impact on the sampling resolution.

Finally, stereo panorama quality control is discussed. Two image quality criteria are taken into account, they are stereo acuity and image composition. Formulas for determining optimum values of R and ω are presented. The determined values ensure the resulting stereo panorama satisfies the specifications of the application requirements regarding these two image quality criteria.

Generalizations in Computer Vision: from Intrinsically Multi-dimensional Signals to Stratified Robot Vision

Gerald Sommer, Michael Felsberg and Bodo Rosenhahn
(Universität Kiel)

Six years ago we presented at the same Dagstuhl seminar a program to overcome some problems of the disciplines contributing to the perception-action cycle. These disciplines are image processing, computer vision, neural computing and robotics. The problems have their origin in the limited representation power of the vector algebra which is the common approach of modelling. We proposed to study the use of geometric algebra or Clifford algebra because of their rich representation structure.

In the present talk we consider two problems of different fields and the way to overcome them. The first problem is the well-known fact that linear op-

erators used in image processing do not respond correctly to intrinsically two-dimensional signals. Although these structures are the most interesting ones, the commonly applied technique of image analysis is incomplete with that respect. We identify the reason for that fact in the limited capabilities of representing the relevant symmetries of such structures in the domain of complex numbers. Especially we want to use the concepts of the analytic signal and of the Hilbert transform also in its generalized versions for intrinsically two-dimensional signals embedded in two dimensional domains. The way to find out the necessary generalizations is rooted in multi-dimensional generalizations of the complex analysis to Clifford analysis. In two dimensions follows from the fundamental solution of the Laplace equation the Riesz transform as isotropically generalized Hilbert transform and the monogenic signal as generalization of the analytic signal. The monogenic signal as local spectral signal representation is adequate for embedding intrinsic 1D-signals into 2D-domains. From that representation follows besides local magnitude and local phase also the local orientation of the signal. Thus, it is not necessary to apply steerable filters because, as a consequence of the embedding, orientation follows immediately. Besides, the embedding gives an additional degree of freedom which is the scale. A Poisson scale space as alternative to the Gaussian scale space is a further important result of the embedding.

The Riesz transform is based on using spherical harmonics of first order. The extension of the approach to intrinsically 2D-signals embedded in 2D-domains requires to use in addition at least spherical harmonics up to order three. Based on such approach a structure multivector can be derived from a set of convolutions with real-valued operators. From the seven components of that structure multivector a rich set of local features can be computed which cover both intrinsically 1D- and 2D- signal properties and which enables to identify the intrinsic dimension.

The second topic is dedicated to the well-known problem of 2D-3D pose estimation. We propose to embed the problem into a conformal geometric algebra. In that framework the Euclidean space, the projective space and the conformal space are stratified with respect to their geometric algebras. All these spaces will be used in the context of the proposed pose estimation algorithm. This approach is using constraint equations between 3D-entities of an object model and projective reconstructed 3D-entities resulting from image features. A further advantage of using the conformal embedding rests on the fact that the basic entities of the conformal space are spheres. From spheres other entities can be derived as points, lines, planes and circles. All these entities can be simultaneously used in a system of constraint equations. Although in conformal geometric algebra the Euclidean transformation is linear as spinor product, we prefer the estimation of the motion parameters not with respect to the Lie group but with respect to the Lie algebra. This method is called the twist approach. It results in nearly video-real-time capability. We discuss further advantages of using the mentioned algebraic embedding of the task of pose estimation.

Bias Correction in Photometric Stereo using Control Points

I. Horowitz and Nahum Kiryati
(Faculty of Engineering, Tel Aviv University)

Photometric stereo is capable of high quality reconstruction of fine shape details. It is however prone to bias, caused by systematic error build-up, due to imperfections in the light sources or in their calibration.

We explore possibilities for correcting the bias, using sparse control points of known 3D locations. Control points can easily be obtained via triangulation, either by projecting a laser pattern or by adding a camera and registering a small number of landmarks. Previous mathematical approaches to the incorporation of control points as constraints in the computation of shape from normal directions lead to inadequate results.

We propose two methods for bias correction using control points. One is based on constrained weighted least squares extension to shape from needle-diagram computation. The other adds an interpolation surface to the reconstructed shape. Experimental results demonstrate significant bias reduction, allowing high reconstruction quality even in the presence of severe setup calibration errors.

Depth from Optical Flow

John L. Barron with Jacky Ngai (University of Western Ontario)
Hagen Spies (Phytosphere, Research Center Jülich)

We present a Kalman filter framework for recovering depth from the time-varying optical flow fields generated by a camera translating and rotating over a scene by a known amount. Synthetic data made from ray traced cubical, cylindrical and spherical primitive objects are used in the optical flow calculation and allow a quantitative error analysis of the recovered depth. A comparison with Hung and Ho's method (PAMI 1999) is presented.

New Methods for Testing Linear Separability

Mohamed Tajine and David Elizondo
(Pole API, LSIIT - UMR 7005 CNRS, Illkirch)

If two sets are linearly separable (LS), then there exists a single layer perceptron feedforward neural network that classifies them. We propose three methods for testing linear separability. The first method is based on the notion of convex hull, the second on a halting upper bound for the perceptron learning algorithm, and the third one, called the class of linear separability algorithm, is based on the characterization of the set of points by which passes a hyperplane that

linearly separates two classes.

We also study the treatment of non linearly separable (NLS) classification problems. We propose two solutions for handling NLS problems. The first solution is an approximation of a NLS classification problem by means of a LS one. We prove that the problem of finding the best approximation is NP-complete. The second solution is a NLS to LS transformation algorithm which can be used for building a multilayer feedforward neural networks called the Recursive Deterministic Perceptron (RDP). This neural network model can be used for solving any two class classification problem even if the two classes are NLS. We prove that the construction of an RDP linearly separating two classes X and Y can be done in less than n steps where n is the cardinal of the set $X \cup Y$.

The geometrical halting upper bound for the perceptron algorithm and the class of linear separability algorithm are the principal results of this paper.

Discrepancy Problems in Data Rounding Related to Digital Halftoning

**Kunihiko Sadakane, Nadia Takki Chebihi and Takeshi Tokuyama
(Tohoku University)**

Digital halftoning is a problem of computing a binary image approximating an input gray image. Although human vision is the traditionally authorized system to evaluate the quality of halftoning, it is desired to establish an automatic evaluation system. We have tested an evaluation system based on discrepancy measures. Our principal approach lead to a range of useful quality measures, that we examine during halftoning process.

Based on experiments using the evaluation system, we have formulated the digital halftoning problem as a matrix rounding problem that converts a real valued two dimensional array into an integral valued array. Since the matrix rounding problem is NP-hard, we designed several heuristic algorithms based on the sequence rounding algorithms, which is its one-dimensional analogue, and succeeded to reduce the discrepancy of the halftoning output.

Polygon Decomposition Based on Straight-line Skeletons

Mirela Tanase and Remco Veltkamp (Utrecht University)

Shape decomposition is of intrinsic interest, because it describes a shape in terms of its constituent parts. In addition, our interest comes from shape-based image retrieval. For such retrieval, one typically wants to be able to perform partial matching: two shapes are considered similar if a part of one shape matches well with a part of another. The problem is that similarity measures for partial matching typically do not obey the triangle inequality.

This is a problem because many (image) database index schemes depend on the triangle inequality to avoid a linear search through the whole database. If however, we decompose the shapes first, and perform matching and indexing on the individual parts, we effectively perform partial matching, and the similarity measure between the parts can obey the triangle inequality.

We base our decomposition on the Straight Line Skeleton. This skeleton has only straight segments (as opposed to the Medial Axis), and results from a wavefront propagating inwards. The question now is how to employ the straight nature of the skeleton, and the events that happen during the wavefront propagation. Our decomposition consists of two phases. In the first phase, all split events are handled, resulting in the first rough decomposition into parts. In the second phase, reflex edge events are used to successively cut off detail from the parts. Preliminary comparison with results from other methods shows that we obtain intuitively appealing decompositions where other methods behave counterintuitively.

Theory of Hyperfigures

Kokichi Sugihara (University of Tokyo)

The conventional Minkowski sum, denoted by $*$, of figures is not invertible in the sense that, given two figures A and B , the equation $A * X = B$ with unknown X is not solvable; sometimes X does not exist, and sometimes X exists but is not unique.

The goal of this talk is to reformulate the definition of the Minkowski sum operation so that the operation is always invertible. For this purpose we consider a closed curve $c(t)$ satisfying the conditions

- (i) $c(t)$ is a mapping of nonnegative reals into the plane,
- (ii) $c(t) = c(t + k)$ for some integer k ,
- (iii) $c(t)$ is continuously differentiable, and
- (iv) $\arg(c'(t)) = 2t\pi$.

Let us denote the set of all such curves by C^+ . For two curves c and d in C^+ , we define a new operation $c * d$ by the pointwise addition of $c(t)$ and $d(t)$ for the same parameter value t . The conventional Minkowski sum of the regions bounded by the outermost parts of c and d coincides with the region bounded by the outermost part of $c * d$. In this sense, our new operation is a generalization of the conventional Minkowski sum. This operation is associative, commutative, and injective. Hence, by the algebraic extension theorem, we can extend the world C^+ to a larger world, say C , in which the operation is always invertible. We call the elements of C *hyperfigures*.

In this extended world, we can take the Minkowski sum and its inverse freely, without worrying about the inconsistency of the result. Also the hyperfigures can be interpreted physically. For example, suppose that we want to compute

the subregion of a flower bed to which we can spread water using a water-spreading vehicle. Conventionally, this computation can be done only when the shape of the flower bed is given. In our new algebraic system, on the other hand, we can compute and represent the water-spread ability of the vehicle by a hyperfigure before we are given the shape of the flower bed.

**Discrete Variational Method
for the Non-Parametric Surface Reconstruction
from Noisy Sample Clouds**

Atsushi Imiya

(National Institute of Informatics and IMIT, Chiba University)

For a tel-surgery by a robot it is required to detect the geometric information of deformable organs in the human body during surgery. If we measure deformable objects using multi-directional camera systems, we obtain a noisy cloud of sample points which distribute around the surface of the deformable object. For the tracking and computation of geometric parameters of deformable objects from the cloud of sample points, we are required to estimate the boundary of an object which lies in the cloud of sample points.

In this paper, we propose a method for estimation of a surface from a cloud of noisy samples. The variational method is employed for the extraction of boundary curves and surfaces for the segmentation of a region and reconstruction. A classical data fitting method based on splines is formulated using the variational method for regression functions. This is a parametric model fitting problem. For data with many outliers, a weighted variational method (M-estimator) is used. This parametric model fitting problem yields smooth solutions as surfaces and boundaries. Variational problems are solved by a scale space method using a diffusion type equation which is derived from the Euler-Lagrange equation of the variational problem. Computational geometry provides combinatorial methods for the recovery of boundary curves and surfaces as polygonal curves and polyhedral surfaces, respectively. These algorithms are based on Voronoi tessellation, Delaunay triangulation, Gabriel graphs, crust, alpha-shape, and beta-skeleton. The reconstructed curves and surfaces by these methods are piecewise linear. Furthermore, the solutions are sensitive to noise and outliers, since these methods construct polygons and polyhedrons using all sample points.

Recently, Osher and his colleagues proposed a non-parametric variational method based on a distance metric. This method requires to solve a partial differential equation numerically.

In this paper, using the discrete formulation, we also propose a non-parametric variational method for the reconstruction of shapes. The proposed method is stable against noise and outliers because the method extracts a medial axis and a medial surface from partitions of the sample points. The method simulta-

neously determines the partitions of a sample cloud and the medial surface of each partition, since the method is an extension of the principal curve extraction. Furthermore, the method extracts closed surfaces and curves although the original principal curve analysis extracts only open planar curves.

The criterion is also valid for a k -dimensional manifold in an n -dimensional Euclidean space, where k is less than $n - 1$. Our method automatically learns the topology of the collection of sample points, first assuming it topology as an open surface, second evaluating a condition, and third detecting the correct topology of a sample cloud which expresses the geometrical concept.

Inverse Quantization in 2D and 3D

Akihiko Torii and Atsushi Imiya (Chiba University)

We propose an inverse quantization method for 2D and 3D binary data. The performance of computers is progressing, and the bandwidth of networks is getting wider. So we can easily exchange data by downloading or uploading through the Web or by searching from a database. However, as bandwidth and resources are limited, efficient data compressing and approximate reconstructing are still of importance.

Our algorithm of reconstruction has a common mathematical method, which is the deformation of control points using curvature flow and an estimation of a C^2 boundary.

The cubic B-spline curve model is appropriate to our problem. Many studies have been reported already on Cubic B-spline curves in image processing. One of the major subjects is boundary estimation from noise, discontinuities, etc. The other is deformation of control points to fit a cubic B-spline curve optimally. The deformation of control points is appropriate to our method. But we propose a total reconstruction algorithm from compressed data.

Our goal is to build a total reconstruction method for 3D binary data. So we first start from proposing an algorithm for boundary extraction and reconstruction for 2D images. Next, we apply the reconstruction algorithm for 3D closed and open curves with some extensions. Finally, we extend the algorithm for boundary extraction and reconstruction to 3D objects and terrain data such as obtained from GIS (Geographic Information System).

Reestablishing Consistency of Uncertain Geometric Relations Detected in Digital Images

Peter Veelaert (Hogent, Ghent)

We explain how uncertain geometry can be used to determine which geometric relations exist between line segments in a digital image. In uncertain geometry

we associate with each point or pixel an uncertainty region. Then a set of points is called straight if there is a Euclidean straight line crossing all the uncertainty regions of the points in the set. Parallelism, collinearity and concurrency are defined in a similar way. For example, two point sets are called parallel if we can find two parallel Euclidean straight lines, such that for each set, one of the Euclidean lines crosses all the uncertainty regions associated with the set. Thus uncertain geometry inherits all its definitions in a straightforward way from Euclidean geometry. An apparent drawback of this approach is that it leads to inconsistencies, due to the introduction of uncertainty. For example, in uncertain geometry parallelism is no longer an equivalence relation. Geometric consistency can be partially restored, however, by grouping processes which are based on minimum clique coverings of graphs that represent the initially derived relations between line segments. In this work, we show that grouping restores consistency only up to a certain level. For example, the geometric relations derived after grouping may still violate well-known geometric theorems such as Desargues's or Pappus's Theorem. We therefore propose a technique to restore global consistency. This technique is based on gradual reductions of the uncertainty regions of line parameters during the grouping process.

Fan Clouds versus Triangular Meshes

Hartmut Prautzsch (Universität Karlsruhe)

To reconstruct a surface, it is common to digitize it, i.e., to determine a sufficiently fine set of points on the surface. Such point sets, or point clouds, are often triangulated since standard visualization and surface editing tools are based on triangular meshes.

In this talk, it is shown that it is often simpler to use fan clouds instead of triangular meshes. Fan clouds consist of local triangulations, i.e., one triangle fan for every point of a point cloud. The advantages are that fan clouds can be computed faster than triangular meshes, and they are better suited for local editing and local refinement.

Potential drawbacks are that in general fan clouds have approximately three times as many triangles than triangular meshes, and that they do not provide continuous surface representations. It is shown how to overcome these potential drawbacks.