


Mathematics and Visualization



Lars Linsen Bernd Hamann
Hans Hagen Hans-Christian Hege *Editors*

Visualization in Medicine and Life Sciences II

Progress and New Challenges

 Springer

Mathematics and Visualization

Series Editors

Gerald Farin

Hans-Christian Hege

David Hoffman

Christopher R. Johnson

Konrad Polthier

Martin Rumpf

For further volumes:

<http://www.springer.com/series/4562>

Lars Linsen
Hans Hagen
Bernd Hamann
Hans-Christian Hege
Editors

Visualization in Medicine and Life Sciences II

Progress and New Challenges

With 132 Figures, 116 in color

 Springer

Editors

Lars Linsen
Jacobs University, Bremen
School of Engineering and Science
Postfach 75 05 61
28725 Bremen
Germany
l.linsen@jacobs-university.de

Bernd Hamann
Institute for Data Analysis and Visualization
Department of Computer Science
University of California, Davis
One Shields Avenue
Davis, CA 95616-8562
U.S.A.
hamann@cs.ucdavis.edu

Hans Hagen
Universität Kaiserslautern
FB Informatik
Postfach 30 49
67653 Kaiserslautern
Germany
hagen@informatik.uni-kl.de

Hans-Christian Hege
Zuse Institute Berlin (ZIB)
Visualization and Data Analysis
Takustr. 7
14195 Berlin
Germany
hege@zib.de

ISBN 978-3-642-21607-7 e-ISBN 978-3-642-21608-4
DOI 10.1007/978-3-642-21608-4
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2007935102

© Springer-Verlag Berlin Heidelberg 2012, Corrected printing 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

Medicine has for a long time been a major driver for the development of data processing and visualization techniques. Many exciting and challenging visualization problems are continually arising in medicine as a consequence of our ability to generate increasingly large and complicated data (image data, simulated data) that require us to devise effective tools for meaningful interpretation and utilization in medical practice. The first VMLS workshop, which led to the book entitled “Visualization in Medicine and Life Sciences (VMLS),” was driven by the fact that emerging technologies in the life sciences had produced significant data visualization challenges. One interesting question was: Can medical data visualization approaches be devised and/or improved to meet these challenges with the promise of ultimately being adopted by medical experts.

Life sciences are understood by us in a broad sense, including animal and human biology, biochemistry, bioinformatics, biomathematics, food sciences, environmental sciences, and pharmacology. Different data acquisition technologies lead to different types of data, including both spatial and non-spatial data. The aim of the second international VMLS workshop was to document and discuss the progress that had been made since the first workshop and to explore what novel solution approaches for data processing and visualization had been developed and what new challenges had come up.

Internationally leading experts from the visualization and driving medical application areas came together for this second workshop held in Bremerhaven, Germany, in July 2009. Research and survey papers were solicited and peer-reviewed, ultimately leading to the collection of papers included in this book.

The research topics covered by the papers in this book deal with these themes:

- Feature Extraction
- Classification
- Volumes and Shapes
- Tensor Visualization
- Visualizing Genes, Proteins, and Molecules

The workshop was supported, in part, by the Deutsche Forschungsgemeinschaft (DFG). We would like to thank Diana Babiac for her help with compiling this book.

Bremen, Germany
Kaiserslautern, Germany
Davis, California, U.S.A.
Berlin, Germany

Lars Linsen
Hans Hagen
Bernd Hamann
Hans-Christian Hege

Contents

Part I Feature Extraction

Discrete Distortion for 3D Data Analysis	3
Leila De Floriani, Federico Iuricich, Paola Magillo, Mohammed Mostefa Mesmoudi, and Kenneth Weiss	

Interactive Visualization—A Key Prerequisite for Reconstruction and Analysis of Anatomically Realistic Neural Networks	27
Vincent J. Dercksen, Marcel Oberlaender, Bert Sakmann, and Hans-Christian Hege	

MRI-Based Visualisation and Quantification of Rheumatoid and Psoriatic Arthritis of the Knee	45
Ben Donlon, Douglas Veale, Patrick Brennan, Robert Gibney, Hamish Carr, Louise Rainford, ChinTeck Ng, Eliza Pontifex, Jonathan McNulty, Oliver FitzGerald, and John Ryan	

An Application for the Visualization and Quantification of HIV-Associated Lipodystrophy from Magnetic Resonance Imaging Datasets	61
Tadhg O’Sullivan, Patrick Brennan, Peter Doran, Paddy Mallon, Stephen J. Eustace, Eoin Kavannagh, Allison Mcgee, Louise Rainford, and John Ryan	

Part II Classification

Semi-Automatic Rough Classification of Multichannel Medical Imaging Data	71
Ahmed Elmoasry, Mohamed Sadek Maswadah, and Lars Linsen	

An Evaluation of Peak Finding for DVR Classification of Biological Data	91
Aaron Knoll, Rolf Westerteiger, and Hans Hagen	
Part III Volumes and Shapes	
Vessel Visualization with Volume Rendering	109
Christoph Kubisch, Sylvia Glaßer, Mathias Neugebauer, and Bernhard Preim	
Efficient Selection of Representative Views and Navigation Paths for Volume Data Exploration	133
Eva Monclús, Pere-Pau Vázquez, and Isabel Navazo	
Feature Preserving Smoothing of Shapes Using Saliency Skeletons	153
Alexandru Telea	
Part IV Tensor Visualization	
Enhanced DTI Tracking with Adaptive Tensor Interpolation	173
Alessandro Crippa, Andrei C. Jalba, and Jos B.T.M. Roerdink	
Image-Space Tensor Field Visualization Using a LIC-like Method	191
Sebastian Eichelbaum, Mario Hlawitschka, Bernd Hamann, and Geric Scheuermann	
Towards a High-quality Visualization of Higher-order Reynold's Glyphs for Diffusion Tensor Imaging	209
Mario Hlawitschka, Younis Hijazi, Aaron Knoll, and Bernd Hamann	
Part V Visualizing Genes, Proteins, and Molecules	
VENLO: Interactive Visual Exploration of Aligned Biological Networks and Their Evolution	229
Steffen Brasch, Georg Fuellen, and Lars Linsen	
Embedding Biomolecular Information in a Scene Graph System	249
Andreas Halm, Eva Eggeling, and Dieter W. Fellner	
Linking Advanced Visualization and MATLAB for the Analysis of 3D Gene Expression Data	265
Oliver Rübél, Soile V.E. Keränen, Mark Biggin, David W. Knowles, Gunther H. Weber, Hans Hagen, Bernd Hamann, and E. Wes Bethel	
Index	285

Part I

Feature Extraction

Discrete Distortion for 3D Data Analysis

Leila De Floriani, Federico Iuricich, Paola Magillo, Mohammed Mostefa Mesmoudi, and Kenneth Weiss

Abstract We investigate a morphological approach to the analysis and understanding of three-dimensional scalar fields, and we consider applications to 3D medical and molecular images as examples. We consider a discrete model of the scalar field obtained by discretizing its 3D domain into a tetrahedral mesh. In particular, our meshes correspond to approximations at uniform or variable resolution extracted from a multi-resolution model of the 3D scalar field, that we call a *hierarchy of diamonds*. We analyze the images based on the concept of *discrete distortion*, that we have introduced in [26], and on segmentations based on Morse theory. Discrete distortion is defined by considering the graph of the discrete 3D field, which is a tetrahedral hypersurface in R^4 , and measuring the distortion of the transformation which maps the tetrahedral mesh discretizing the scalar field domain into the mesh representing its graph in R^4 . We describe a segmentation algorithm to produce Morse decompositions of a 3D scalar field which uses a watershed approach and we apply it to 3D images by using as scalar field both intensity and discrete distortion. We present experimental results by considering the influence of resolution on distortion computation. In particular, we show that the salient features of the distortion field appear prominently in lower resolution approximations to the dataset.

L. De Floriani (✉) · F. Iuricich · P. Magillo · M.M. Mesmoudi
Department of Computer Science and Information Science (DISI), University of Genova, Italy,
Via Dodecaneso 35, 16146 Genova (Italy)
e-mail: defflo@disi.unige.it; federico.iuricich@gmail.com; magillo@disi.unige.it;
mmesmoudi@creteil.fr

K. Weiss
Department of Computer Science, University of Maryland, College Park, MD 20742 (USA)
e-mail: kweiss@cs.umd.edu

M.M. Mesmoudi
Labo. Pure and Appl. Math. UMAB, University of Mostaganem, Route Belhacel, 27000
Mostaganem (Algeria)

1 Introduction

We consider a three-dimensional scalar field which is defined by a collection of function values, each given at a point in a 3D domain. Examples of 3D scalar fields of interest in biomedical applications are 3D images, where the intensity at each voxel defines the scalar field. A scalar field is known at a finite set of points in 3D space, and a digital model of the field is constructed based on such points. Models of 3D scalar fields are based on voxels, or on tetrahedral meshes with vertices at the points in the 3D space at which the field is known. In both cases, such models tend to be verbose and may not be immediately useful to understand the behavior of the field.

Here, we consider tetrahedral meshes extracted from a multi-resolution representation of 3D images provided by a regular tetrahedral hierarchy. We have developed an efficient representation of a regular tetrahedral hierarchy, called a hierarchy of diamonds, as discussed in [43].

The aim of morphological analysis is to provide a tool for understanding the structure of a scalar field through structural representations of the field so that its basic features can be easily recognized. Here, we use the notion of discrete distortion to support morphological analysis. In [26] we have introduced a discrete approach to curvature for three-dimensional tetrahedralized shapes embedded in 4D space, that we called *discrete distortion*. If we consider 3D scalar fields, we can view the values of the field as constraints on the vertices of a tetrahedral mesh. From this perspective, the values induce a distortion of the geometry of the mesh, seen as a hypersurface representing the graph of the scalar field in R^4 . As for surface curvature, discrete distortion highlights the local curvature of the constrained shape (the graph of the 3D scalar field) which cannot be perceived in the three-dimensional domain. As curvature gives interesting insights in terrain analysis, we show that distortion provides additional information to analyze the behavior of the intensity field. A null distortion value highlights a linear behavior of the intensity field, while a constant distortion corresponds to a uniform non-linear behavior. We observe that directions in which distortion changes indicate interesting directions in which the intensity field varies its growth speed.

One way to perform morphological analysis is to automatically decompose the domain of the field into meaningful parts in order to support understanding and semantic annotation. Segmentation has been the basic tool to support reasoning on terrains and 3D shapes. Here, we propose segmentations for a 3D image based the intensity value or on discrete distortion, in a similar way as done for terrains where segmentations are computed based on elevations and/or on curvature values.

The segmentation of a scalar field is performed based on its critical points, and the steepest directions through which the scalar field increases or decreases. This leads to two dual decompositions. The *stable Morse decomposition* associates a 3D cell with each local minimum of the field, and two adjacent 3D cells touch at ridge surfaces (i.e., surfaces where the field decreases on both sides). The *unstable Morse decomposition* symmetrically associates a 3D cell to each local maximum, and the

boundaries of the 3D cells are at valley surfaces. Here, we present an algorithm for computing Morse decompositions based on a watershed approach and we compute Morse decompositions based on the intensity field and on a new field induced by discrete distortion. This approach reveals relevant features, different from those that are generally extracted by studying the behavior of the gradient field (i.e., the critical points of the intensity field). For example, the extrema of distortion correspond to locations in which the intensity field has abrupt variations, which might not be perceptible from the intensity values.

We apply our approach to the analysis of the morphology of scalar fields through examples on synthetic, biological and medical datasets. We show color-coded visualization of the fields based on the intensity field and of distortion. We study the interaction between the resolution of the tetrahedral mesh approximating the field and the distortion values, showing that we can reasonably approximate the 3D image at fairly low resolutions. Finally, we show results on Morse segmentations based on the intensity and on the distortion values and we compare them.

The remainder of this paper is organized as follows. In Sect. 2, we provide background notions on concentrated curvature, and on Morse theory and Morse complexes. In Sect. 3, we review some related work. In Sect. 4, we briefly describe a mesh-based multi-resolution model that we use for representing the 3D image. In Sect. 5, we discuss the notion of discrete distortion for a tetrahedralized shape representing the graph of a 3D scalar field, and we present some of its properties. In Sect. 6, we present an algorithm to segment tetrahedral meshes endowed with discrete scalar fields, and produce Morse decompositions. In Sect. 7, we present experimental results on medical data set and we discuss the results. Finally, in Sect. 8, we draw some concluding remarks and discuss on-going and future work.

2 Background Notions

In this Section, we discuss first concentrated curvature, since the notion of discrete distortion generalizes concentrated curvature to hypersurfaces in R^4 . Then, we briefly review some notions from Morse theory which is the basis for defining morphological segmentations for a scalar field.

2.1 Concentrated Curvature

Concentrated curvature is the discrete counterpart of Gaussian curvature for triangulated surfaces [1, 25, 40], and was introduced by Aleksandrov for 2D scalar fields represented as triangle meshes [1]. Given a triangulated surface in R^3 and a vertex p in the interior of the corresponding triangle mesh, the local neighborhood of p is the union of the angular sectors of the triangles incident at p . The *total*




Fig. 1 In (a), surface of a molecule with associated concentrated curvature represented in a rainbow color scale (blue = low values, red = high values). In (b), Morse decomposition of the surface based on concentrated curvature

angle Θ_p at p is given by the sum of the angles at p of all the triangles incident in p . The *concentrated curvature* at a vertex p is defined as $K(p) = 2\pi - \Theta_p$, when p is an internal vertex [40]. If the sum of the angles of all these sectors is not equal to 2π , p is called a *singular conical point*. When the surface is defined by a scalar field, the concentrated curvature for boundary points can be defined as the defect between the angle at the vertex on the surface and its corresponding angle on the xy -projection domain.


Note that concentrated curvature involves only the sum of angles of the triangles incident at a vertex and does not take into account their geometric position. *Planar points* have null concentrated curvature. *Saddle points* have negative concentrated curvature, while *convex and concave points* have positive concentrated curvature. Figure 1a illustrates the behavior of concentrated curvature on the surface of a molecule: red areas correspond to high curvature, and blue areas to low curvature. The segmentation in Fig. 1b represents a Morse segmentation of the surface based on concentrated curvature.

2.2 Morse Theory and Morse Complexes

Let $f(x, y, z)$ be a scalar field defined on a domain D of R^3 , and let function f be continuous and smooth in D . A point p of D is a *critical point* of f if the gradient of f at p is null. Points that are not critical are called *regular points*. A critical point p is *degenerate* if the Hessian matrix of the second partial derivatives of f at p is not singular. A scalar field f is a *Morse function* if and only if all its critical points are not degenerate. Morse [27] showed that the critical points of a Morse function are isolated. The number i of negative eigenvalues of the Hessian matrix is called the *index* of critical point p , and p is called an i -*saddle*. A 0-saddle is a *minimum* and a 3-saddle a *maximum*.

An *integral line* of f is a maximal path which is everywhere tangent to the gradient of f . Each integral line connects two critical points of f , called its *origin* and its *destination*. Integral lines that converge to a critical point p of index i form

Fig. 2 Stable Morse complex for a 2D scalar field. Arrows denote the negative gradient field. Only a window over the domain of the field is shown; the stable cells of the minima are shown in different colors: each stable cell is the region covered by arrows converging to the same minimum



an i -cell, called a *stable* cell of p . Dually, integral lines that originate at p form its *unstable* ($n - i$)-cell. The stable and unstable cells decompose D into *stable* and *unstable Morse complexes*. Figure 2 illustrates the above concepts for the domain of a 2D scalar field.

A Morse function f is called a *Morse-Smale function* if each non-empty intersection of a stable and an unstable cell is transversal. The connected components of the intersection of the stable and unstable cells define a *Morse-Smale complex*. Note that the Morse-Smale complex for f can be obtained by the overlay of its stable and unstable complexes.

3 Related Work

In this Section, we briefly review related work on discrete curvature estimators, on multi-resolution modeling of 3D scalar fields, on algorithms for computing Morse decompositions of 3D scalar fields, and on features of interest in medical imaging.

3.1 Discrete Curvature Estimators

Curvature is an important notion in mathematics that found a great interest in the last century. Curvature is also used to study the local geometry and topology of surfaces from the metric point of view. With the development of discrete geometry, many authors tried to define a discrete counterpart of curvature based on the properties observed in the continuum [12, 18, 37, 38]. There is a rich literature dealing with the problem of defining and computing discrete curvature estimators for triangle meshes, and more recently for tetrahedral meshes (see [12, 18, 37, 38] for a survey). Concentrated curvature [1, 40] is a simple and efficient method to define a discrete curvature, as discussed in Sect. 2.1. Dyn et al. discuss how to optimize the triangulation of the boundary of a 3D object based on discrete curvature [9].

In the 3D case, the Ricci tensor is used to define the curvature notion for three-dimensional shapes [2], and, in the discrete case, the Laplace operator is generally used to define a discrete approach to curvature [33].

In the 4D case, curvature is one of the most important mathematical notions on which general relativity is based. Curvature of the space-time gave an important contribution to understand many phenomena in physics (black holes, gravitational lenses, light trajectories, interaction between planets, ...). Based on Aleksandrov's concentrated curvature, Regge introduced a discrete version of curvature for the four dimensional space-time [32].

3.2 *Multi-Resolution Modeling*

A very large class of multiresolution models of volumetric scalar fields is provided by nested meshes, in which all elements are defined by the uniform subdivision of a small set of primitive cells. Examples include octrees formed by cubes [36] and tetrahedral meshes generated by the so-called *Red/Green* tetrahedron refinement [5].

Nested tetrahedral meshes based on the *Longest Edge Bisection (LEB)* operation were originally introduced for domain decomposition in finite element analysis [20, 23, 34], and have since then been applied in many different contexts, including scientific computing [13, 14, 46], surface reconstruction [24] and volume segmentation [21]. A recent survey on nested simplicial meshes based on bisection can be found in [44].

The *LEB* operation is defined by bisecting a tetrahedron t along the plane defined by the midpoint of its longest edge e and the two vertices of t not incident to e . The containment relation among the tetrahedra generated by successive *LEB* operations naturally defines a binary tree, where the two tetrahedra generated by bisecting a *parent* tetrahedron t are the *children* of t . When a full binary tree is stored, this representation can be efficiently encoded as a linear array, and the parent-child relation can be implicitly determined from the array indices [13, 22, 46]. A forest of six such tetrahedral binary trees, whose roots share a common cube diagonal can thus decompose a cubic region of space.

We are often interested in generating crack-free, or *conforming*, tetrahedral meshes, since cracks in the mesh correspond to discontinuities in scalar fields discretized through it. Methods of ensuring continuity have been proposed based on a hierarchical monotonic error metric [29], symbolic *neighbor-finding* operations [20, 22] or an implicit clustering of tetrahedra sharing a common bisection edge into a *diamond* primitive [14, 42]. In this work, we utilize diamonds to extract conforming tetrahedral meshes from the multi-resolution model.

3.3 *Algorithms for Morse Decompositions in 3D*

Most of the algorithms proposed in the literature for extracting an approximation of the Morse complexes in the discrete case have been developed for terrains. The majority of them use a *boundary-based* approach, since they extract the

decomposition by computing the critical points and then tracing the integral lines, or their approximations, starting from saddle points and converging to minima and maxima. Other algorithms use a *region-based* approach, in the sense that they compute an approximation of a Morse decomposition by growing a 2D region defined and started by the minima and the maxima of a Morse function f . Curvature has been applied to the segmentation of 3D shapes and terrains in combination with Morse decompositions (see, for instance, [28]). A comprehensive analysis of techniques for Morse decomposition can be found in [6].

Alternative region-based techniques for computing the stable and unstable Morse decompositions are those based on the discrete watershed transform (see [35] for a survey). For a C^2 -differentiable function f , the watershed transform provides a decomposition of the domain of f into regions of influence of the minima, called *catchment basins*, which are bounded by watershed lines. If f is a Morse function, it can be shown that the catchment basins of the minima of f and the watershed lines correspond to the 2-cells and the 1-cells, respectively, in the stable Morse decomposition of f [6].

Much less work has been done on computing Morse decompositions for 3D scalar fields. In [10], an algorithm for extracting the Morse-Smale decomposition (i.e. the intersection of the stable and unstable Morse decompositions) from a tetrahedral mesh approximating a 3D scalar field is proposed. The algorithm computes the Morse-Smale decomposition by extracting the critical points, then the unstable Morse decomposition and finally the stable cells in pieces inside the unstable cells. The algorithm, while interesting from a theoretical point of view, has a large computation overhead, as discussed in [17]. In [17], a region growing method is proposed to compute the Morse-Smale decomposition inspired by the watershed approach. A procedural approach based on discrete Morse theory is described in [15] which also computes the Morse-Smale complex.

A major issue when computing Morse decompositions for 2D and 3D scalar fields is over-segmentation, which is due to the presence of noise in the data. To this aim, *generalization algorithms* have been developed in order to eliminate less significant features from a Morse or Morse-Smale decomposition, mainly for 2D scalar fields. Generalization is achieved by applying an operation, called *cancellation* of critical points. Cancellations of critical points for a 3D scalar field consist of collapsing a maximum and a 2-saddle into a single maximum, a minimum and 1-saddle into a single minimum, or a 1-saddle and a 2-saddle into either a 1-saddle or a 2-saddle [7, 16].

3.4 Features in Medical Images

Features of interest in medical images (also called *landmarks*) may correspond to points, lines, surfaces or volumes. Many techniques for landmark extraction are based on curvature. In [3, 31, 39] an algorithm that computes a polygonal approximation of the so-called *Gaussian frontier* is described. Points with large curvature

values along a contour are selected following different scale-space Gaussian filters. In [4], isolines of extremal values of mean curvature are selected to segment MRI and CT-scans images. These isolines generally correspond to ravine and ridge lines of the surface. Based on contour detection and curvature, an automatic landmark extraction is proposed in [11]. Several techniques using first and second differentials of 3D operators (that extend the 2D case) are compared on MRI and CT-scan images in [19].

4 A Diamond-Based Multi-Resolution Model

In this Section, we describe a mesh-based multi-resolution model for 3D scalar fields called a *hierarchy of diamonds*, which we use to generate approximate representations of a scalar field at variable and uniform resolutions. We have studied the theory and the properties of hierarchy of diamonds in arbitrary dimensions in [42].

A multi-resolution model M of a shape Σ is typically defined by three components [8]: (a) A coarse base mesh Γ_0 that approximates Σ , (b) a set of modifications U , each of which replaces a set of cells γ_1 with a new set of cells γ_2 sharing the same combinatorial boundary and (c) a *direct dependency relation* R on the modifications U , where a modification u_2 depends on another modification u_1 if u_2 replaces a cell that was introduced in u_1 .

As pointed out in Sect. 3.2, the LEB (longest edge bisection) operation is defined by bisecting a tetrahedron t along the plane defined by the midpoint of its longest edge e and the two vertices of t not incident to e .

A *Hierarchy of Diamonds* Δ is a multi-resolution representation of a regularly sampled 3D scalar field covering a cubic base domain D and generated through longest edge bisection. It is based on clusters of tetrahedra, called *diamonds*, sharing the same bisection edge, that we call the *spine* of the diamond. The diamond whose spine is a diagonal of D defines the base mesh Γ_0 of the model. Each diamond δ corresponds to a modification (γ_1, γ_2) in the multi-resolution model, where γ_1 consists of the tetrahedra of δ , and γ_2 consists of the tetrahedra generated by bisecting the tetrahedra in γ_1 along the spine of δ .

A diamond δ_p is said to be a *parent* of another diamond δ_c if one or more of the tetrahedra in δ_c is generated during the bisection of the tetrahedra in δ_p . This parent-child relation defines the *direct dependency* relation of the model, which can be encoded as a *Directed Acyclic Graph* (DAG) (see Fig. 4(a) for an illustration in 2D).

An explicit encoding of the hierarchy as a directed acyclic graph would require each modification to list the tetrahedra before and after bisection as well as the dependencies among these modifications. However, due to the regularity of the vertex distribution and the subdivision rule, this model generates only three *classes* of diamonds composed of six, four and eight *similar* tetrahedra [23], respectively (see Fig. 3). Furthermore, the diamond classes have three, two and four parents, and six, four and eight children, respectively.




Fig. 3 The three classes of three-dimensional diamonds. The spine (internal green edge) of a diamond of class i is aligned with the diagonal of a $(d - i)$ -cube. A diamond's central vertex v_c (hollow red circle) coincides with the midpoint of its spine

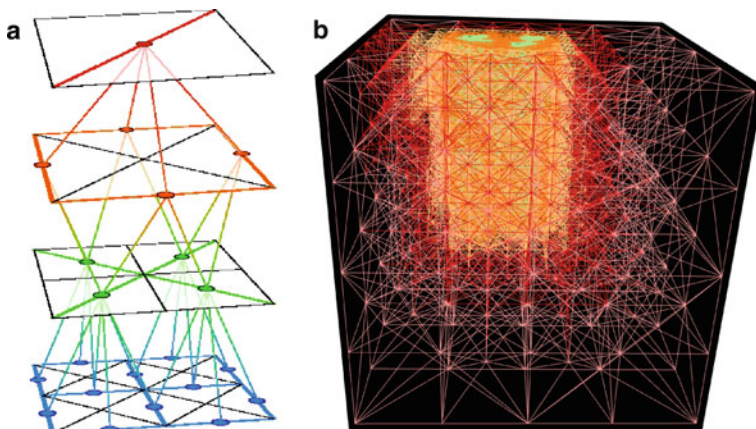


Fig. 4 (a) A hierarchy of diamonds (shown in 2D) is a rooted DAG whose nodes are the diamonds and whose arcs encode the dependency relation among the diamonds. (b) Example variable-resolution tetrahedral mesh extracted from a hierarchy of diamonds. A tetrahedron's color indicates its level of resolution

Thus, diamond hierarchies admit extremely compact encodings of the underlying multi-resolution model which exploit the implicit relationships among the modifications and their dependencies. Each diamond is completely defined by its spine, and all its tetrahedra are split by the diamond's *central vertex*, the unique midpoint of its spine. Thus, diamonds are in one-to-one correspondence with their central vertices, which, in turn, are in one-to-one correspondence with the samples of the dataset. From the coordinates of the central vertex, we use bit manipulations to extract the complete parent-child relations. A hierarchy of diamonds can therefore be encoded as the collection of the central vertices of its diamonds [42], from which all geometric and hierarchical relationships can be implicitly determined [14, 43].

A hierarchy of diamonds Δ is used to efficiently extract variable-resolution tetrahedral meshes Σ approximating a 3D image while satisfying an application-dependent *selection criterion*. The selection criterion can be defined on properties of the domain, such as proximity to a specified region of interest, or on properties of

the range such as its degree of approximation to the underlying dataset. In contrast to octree-based approaches, such tetrahedral meshes have a higher degree of adaptability to the selection criterion [30], and are guaranteed to be free of cracks. Figure 4(b) shows a variable-resolution tetrahedral mesh extracted from a hierarchy of diamonds.

5 Discrete Distortion

In this Section, we briefly review the notion of distortion that we have introduced in [26] and we describe some of its properties. *Discrete distortion* is a generalization of concentrated curvature to tetrahedral meshes in 4D space [26]. In [25], we have investigated discrete distortion for triangulated terrains, and we have shown that it behaves as a discrete counterpart of mean curvature. Here, we discuss the notion of distortion for tetrahedral meshes endowed with a scalar field.

The graphical representation of a scalar field f defined on a tetrahedral mesh Σ is a hypersurface $(\Sigma; f)$ in R^4 , namely, a tetrahedral mesh embedded in R^4 . Hypersurface $(\Sigma; f)$ is generally curved due to the effects of the scalar field values. As for concentrated curvature, one may compare the defect solid angle at the vertices of Σ , when applying the scalar field.

The *distortion* at a vertex p of Σ is defined as the quantity

$$D(p) = 4\pi - \sum_{T_i \in T(p)} S_i, \quad (1)$$

where S_i is the solid angle, after applying the scalar field (i.e., within $(\Sigma; f)$), at vertex p of tetrahedron T_i , and $T(p)$ is the set of all tetrahedra incident at p . A similar formula holds for boundary vertices:

$$D(p) = \sum_{t_i \in t(p)} s_i - \sum_{T_i \in T(p)} S_i, \quad (2)$$

where s_i is the solid angle, within Σ , at vertex p of tetrahedron t_i , and $T(p)$ is the set of all tetrahedra incident at p .

Discrete distortion for 3D scalar fields has similar properties as concentrated curvature for 2D fields. Concentrated curvature gives positive values to locally convex, or concave, areas of the surface, negative values to saddles, and null values to flat areas. Similarly, positive values of distortion correspond to locally convex, or concave, portions of the hypersurface which is the graph of the field. Negative values correspond to saddle and degenerate saddle configurations.

Constant scalar fields are distortion-free (i.e., their distortion is null). This can easily be understood since, for a constant scalar field, mesh $(\Sigma; f)$ is only a translation, in the fourth dimension, of the mesh Σ decomposing the domain of

the field. Hence, the Euclidean geometric structure of the mesh is preserved. More generally, affine scalar fields are distortion-free, since they combine rotations and translations of the whole mesh. Hence, the geometrical structure is not subject to any distortion.

As a consequence, piecewise linear scalar fields are distortion-free at the interior vertices of regions where the field is linear, as they act affinely within such regions. Another relevant property is that distortion is mesh-dependent. This means that the distortion value at a vertex depends on the way in which the neighborhood of such vertex is triangulated.

6 Computing Morse Complexes

We compute a discrete approximation of the unstable and the stable Morse complexes for a 3D scalar field f defined at the vertices of a tetrahedral mesh Σ by extending the watershed approach by simulated immersion developed for 2D images in [41]. We describe only the algorithm for the stable complex since the unstable complex can be built by considering field $-f$.

The watershed algorithm performs the following three steps:

1. Sort the vertices of the mesh by increasing values of field f .
2. Associate all mesh vertices with a local minimum. This is done starting from minima and proceeding based on increasing field values and on increasing distance from already discovered local minima.
3. Assign each tetrahedron to the stable 2-cell of a local minimum, based on the assignments of its vertices.

In the second step we process the vertices of the mesh according to their field values. Let h be the current field value (initially, h is the minimum field value over the mesh). We consider the set H of all the vertices whose field value is equal to h . A priority queue is used to ensure the processing vertices in H in increasing distance from an already assigned vertex. We iteratively pick the first vertex $v \in H$ from the priority queue, and we check if some of its neighboring vertices has already been assigned to a local minimum. If they are all either unassigned, or assigned to the same local minimum, then v is assigned to that local minimum. If two or more neighboring vertices are assigned to different local minima, then v is marked as a watershed vertex. After assigning v , the priorities of the unassigned neighbors of v are updated in the priority queue. The above process is repeated until no more assignments are possible (i.e., the priority queue is empty). Then, for each vertex $w \in H$ that is still unassigned, w is marked as a new local minimum, and all vertices with the same field values equal to h , which are connected to w , are assigned to w . Now, all vertices in H have been assigned to some local minimum (possibly equal to the vertex itself), and the algorithm proceeds with the next field value.

The third step examines each tetrahedron τ and assigns it to a local minimum based on its vertices. If all four vertices of τ are marked as watershed, then τ is

marked as a watershed tetrahedron. Otherwise, among the local minima assigned to the vertices of τ , we choose the one having minimum field value and assign τ to such minimum.

The watershed algorithm described above often produces an over-segmentation of the graph of the scalar field. To overcome this problem, we perform an iterative merging process of the 3D cells of the stable (unstable) Morse complex, which correspond to a topological generalization operator, namely, to the cancellation of a maximum and a 2-saddle in the stable complex, or of a minimum and a 1-saddle in the unstable complex. We consider pairs of adjacent 3D cells in the stable (unstable) Morse complex. We associate with each cell γ_1 a *saliency* measure with respect to any adjacent 3D cell γ_2 , which takes into account: (i) the maximum field difference between the local minimum [maximum] associated with cells γ_1 and γ_2 and the largest [smallest] field value of a point on the common boundary 2-cell of γ_1 and γ_2 ; (ii) the product of the sizes of γ_1 and γ_2 ; (iii) the extent of the common boundary of γ_1 and γ_2 .

Thus, priority is given to merging cells with small differences in field values, with a small area and to cells with long common boundaries. The first measure is the persistence value, as defined in [10, 16].

7 Experimental Results

In this Section, we present some experimental results which show the behavior of discrete distortion as a tool for analysis of 3D images. Because of the large size of current data sets, it is also important to perform accurate analysis on low-resolution representations of the field. Here, we study the influence of mesh resolution on distortion by considering variable-resolution conforming tetrahedral meshes extracted from a hierarchy of diamonds according to a user-defined threshold on the approximation error. In this case, resolution can be coarsened locally in less interesting regions, without affecting the quality of the approximation. Finally, we show and compare segmentations of the 3D images obtained through Morse decompositions of the intensity and of the distortion fields. To this aim, we present results on a synthetic data set in which the intensity field is defined by an analytic function and on two real data sets.

7.1 3D Datasets and Distortion

Our first example is a synthetic dataset defined over a regularly sampled domain of 65^3 vertices. The intensity field is obtained by sampling the analytic function $f(x, y, z) = \sin(x) + \sin(y) + \sin(z)$. We show here the tetrahedral mesh at full resolution extracted from a hierarchy of diamonds built on such data set. It is composed of 275 K vertices and 1.57 million tetrahedra. The relationship between




Fig. 5 Intensity field (left) and distortion field (right) for synthetic data set sampling function $f(x, y, z) = \sin(x) + \sin(y) + \sin(z)$ over a 65^3 grid

the intensity field and the induced distortion field over this domain is illustrated in Fig. 5 along the boundary of the cubic domain using a blue-red color scale to indicate the low and high scalar and distortion values.

The second dataset, called Neghip, is a simulation of the spatial probability distribution of electrons in a high potential protein molecule. The knowledge of electron distribution within such molecules is important in pharmacology to understand the interactions between molecules and an organism. The inhibition of some protein molecules can reduce complications in diseases such as cataracts and neuropathies for diabetic subjects. The understanding of the catalytic mechanism and the electrostatic potential of the molecule plays a relevant role here. It may help to study, at the atomic scale, the transfer of electrons and protons in complex biological processes such as oxidation/reduction in relation to metallic ions by considering the reaction between hemoglobin (containing iron ions) and the oxygen molecule.

In Fig. 6, we show the intensity field and distortion field for the tetrahedral mesh extracted from the Neghip hierarchy at variable resolution corresponding to 0% approximation error. The mesh has 129 K vertices and 728 K tetrahedra. The range of colors used for visualization goes from blue for low values to red for high values, with gray indicating mean values. Discrete distortion highlights the growth behavior of the density scalar field, which is maximal around the atoms. We see that the density field grows quickly around atoms within small regions and then stabilizes its growth. Distortion becomes nearly constant in such case. We observe also that, within regions where the electron density has low values, many small regions have high distortion values. This indicates changes in the electron density and may be due to the interference between adjacent atoms or to some artifacts in the processing of the data. Regions in blue (for distortion) indicate that the scalar field grows differently in different directions. This corresponds to saddle regions where the convexity of the electron density field changes.

The third dataset, called CTA-Brain, is a CTA-scan of a human brain with an aneurysm. Computed Tomographic Angiography (CTA) is a minimally invasive technique that uses imaging technologies (e.g., X-rays) to explore the structure




Fig. 6 Twelve equally spaced slices (along the z-axis) of the intensity field (left), and the distortion field (right) of the Neghip dataset at 0% error. The colors of the distortion field are scaled to highlight the extreme values

of vessels and tissues. A contrast agent is generally used to produce clear images. The original dataset has $512 \times 512 \times 120$ vertices and measure the intensity of the contrast agent. To show the behavior of the intensity field and of distortion, we have extracted a variable-resolution mesh from the diamond hierarchy, which has 1.74 million vertices and 9.52 million tetrahedra.

Figure 7 illustrates the dataset, where the scalar field corresponds to the intensity of the contrast agent, and its distortion, through equally spaced horizontal slices.




Fig. 7 Seven equally spaced slices (along the z-axis) of the CTA-brain dataset at 10% error illustrating the scalar field (left), and the distortion field (right). The colors of the distortion field are scaled to highlight the extreme values

The geometric structure of the scanned region is well represented by distortion. We see that most regions have gray or light blue color, which indicates a uniform distribution of the contrast agent within the brain. The regions with high distortion correspond to changes in the intensity of the contrast product.

7.2 *Distortion and Mesh Resolution*

We now demonstrate the validity of distortion analysis on lower resolution approximations by considering the distribution of distortion values over a set of extracted meshes with increasingly fine resolution. For brevity, we show results only on the two real data sets, Neghip and CT-Brain.

In the first case, we generate a diamond hierarchy Δ_H based on the intensity values of the 64^3 Neghip dataset, which contains 262 K vertices. The error of a diamond δ is computed as the maximum difference between the intensity values of all grid points within the domain of δ , and the value obtained by linear interpolation over the vertices of δ 's tetrahedra. We extract a series of meshes




Fig. 8 Cumulative distribution functions of distortion values (horizontal axis) over increasingly fine meshes extracted from the Neghip dataset

Σ_{ϵ_i} of uniform approximation error ϵ_i from Δ_H , using threshold values of $\epsilon_i \in \{30\%, 10\%, 5\%, 2\%, 1\%, 0\%\}$ of the total error and then evaluate the distortion of the vertices of these meshes.

Figure 8 shows the *Cumulative Distribution Function (CDF)* of the discrete distortion (horizontal axis) of the vertices of each mesh. The sharp spike in the CDF of all datasets around a null distortion value indicates that the vast majority of vertices have (nearly) null distortion. As the resolution increases, this spike becomes steeper, indicating that the increased resolution is distributed among regions with nearly null distortion. Thus, the distortion is concentrated in relatively few vertices within the mesh, and appears prominently in lower resolution approximations. For example, when $\epsilon = 0$, more than 94% of the 129 K vertices in $\Sigma_{0\%}$ have distortion $D(v) \leq |1|$, and for $\epsilon = 2\%$, more than 83% of the 33 K vertices in $\Sigma_{2\%}$ have distortion $D(v) \leq |1|$.

Similarly, Fig. 9 shows the CDF of meshes Σ_{ϵ_i} using threshold values of $\epsilon_i \in \{99\%, 75\%, 50\%, 30\%, 10\%, 5\%\}$ extracted from the CTA-Brain dataset. These meshes illustrate the same general trend as the Neghip approximations, although they are a bit noisier since they are scanned images.

We have obtained similar results for several other datasets in other application domains [45]. These experiments indicate that we can obtain a fairly accurate understanding of the image via its discrete distortion even at lower resolutions, without the need to compute the distortion on the full image.

7.3 Morse Decompositions

In this Subsection, we show Morse decompositions of the synthetic and real data sets computed using the intensity and the distortion fields.




Fig. 9 Cumulative distribution functions of distortion values (horizontal axis) over increasingly fine meshes extracted from the CTA Brain dataset

Let us consider the distribution of the intensity and distortion values for the synthetic data set shown in Fig. 5. Figure 10 shows the stable and unstable Morse decompositions computed based on the intensity and on the distortion fields. It is clear how the distribution of the intensity and of the distortion values influences the corresponding segmentations. Both stable and unstable Morse decompositions obtained from the intensity field consists of 1,331 cells and have a regular structure. The stable decomposition obtained from the distortion field consists of 12,972 cells, while the unstable one consists of 3,738 cells. The decomposition pattern in the stable and unstable distortion-based complex varies in different portions of the mesh. This is due to the function sampling that is different from its period.

Figure 11 shows Morse decompositions built from the full-resolution tetrahedral mesh discretizing the Neghip dataset. We thresholded the visualization along an isovalue to better illustrate the structure of the molecules. The stable and unstable Morse decompositions obtained from the intensity field consist of 104 cells and of 41 cells, respectively. The stable and unstable Morse decompositions obtained from the distortion field consist of 3,654 stable cells and 23,334 cells, respectively. Some components of the unstable decomposition represent the location of atoms (i.e. maxima of the density) and the proper space in which electrons revolve around. Due to the interference of electron density of adjacent atoms, some components are created and correspond to some maxima of the density field. These components do not properly contain atoms.

Figure 12 illustrates the intensity field, the corresponding distortion values and the segmentations obtained from a uniform resolution mesh $\Sigma_{10\%}$ extracted from the CTA-Brain dataset (see also a view as set of slices in Fig. 7). The decomposition obtained from the intensity field consists of 37,631 stable cells and of 23,835 unstable cells, while the decomposition obtained from the distortion one consists

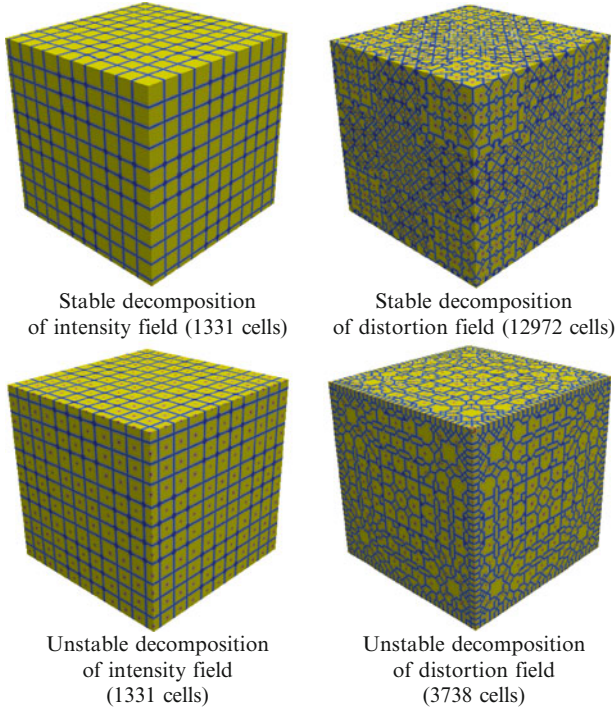


Fig. 10 Morse decompositions for the synthetic data set defined by intensity function $f(x, y, z) = \sin(x) + \sin(y) + \sin(z)$. Minima (stable) or maxima (unstable) vertices are colored in red, vertices on the boundary of several regions in blue and vertices within a region in yellow

of 136,641 stable cells and 128,687 unstable cells. Figure 13 shows the largest segments from the segmentations. Observe that, while the unstable regions are more structured and follow the field values, the stable regions are much more influenced by the boundary and by the less relevant regions of the original scalar field. The former therefore seem to provide a more meaningful decomposition. The large number of cells in the unstable decomposition computed on the basis of distortion is due to the fact that there is a large number of small areas in which the concentration of the contrast agent changes abruptly (i.e., distortion has a maximum).

8 Concluding Remarks

We have presented an innovative approach to the analysis of 3D images based on the notion of discrete distortion, which generalizes discrete curvature to triangulated hypersurfaces in 4D space, and on Morse decomposition.

We have proposed the use of a multi-resolution model based on clusters of tetrahedra, called diamonds, which enables the analysis of a 3D image through

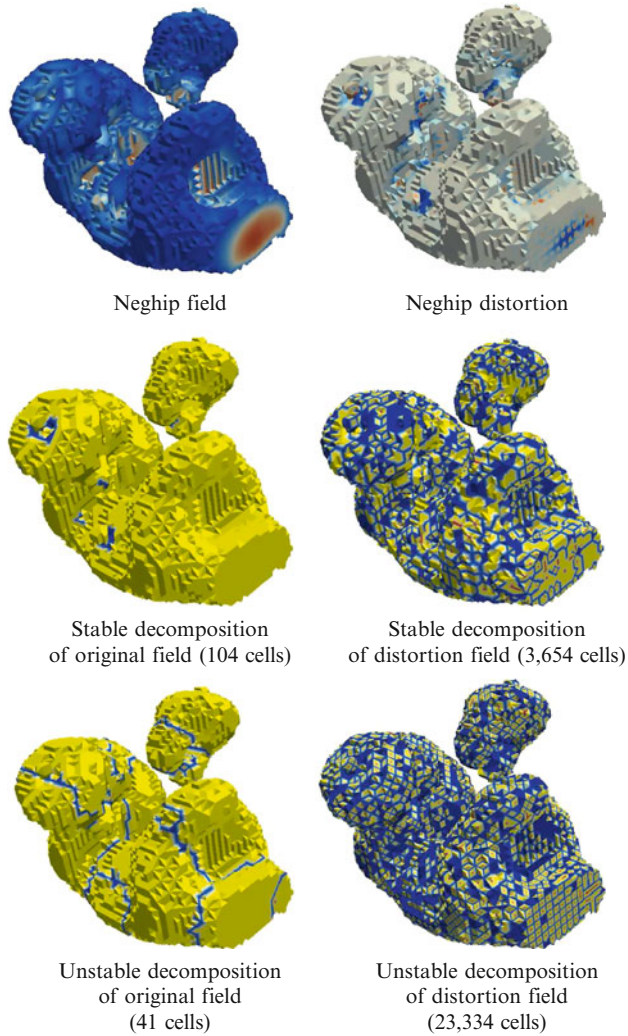


Fig. 11 Original field and distortion field, and segmentations, for variable resolution Neghip data set at 0% approximation error. Segmentations are shown with minima (stable) or maxima (unstable) vertices in red, vertices on the boundary of more than one region in blue and vertices within a region in yellow

crack-free approximations encoded as tetrahedral meshes. One important aspect of using mesh-based multi-resolution models is that the image can be analyzed by using much fewer samples than in the full image. This facilitates our analysis of large 3D volume datasets by using significantly fewer resources. The other aspect that we have shown through our experiments is the utility of discrete distortion in analyzing approximated images, thus giving good insights about the field behavior already at low resolutions.