

Multiresolution Analysis of 3D Images Based on Discrete Distortion

Kenneth Weiss¹, Leila De Floriani², Mohammed Mostefa Mesmoudi²

¹ Department of Computer Science, University of Maryland, College Park.

² Department of Computer Science, University of Genova, Italy.

Abstract

We consider a model of a 3D image obtained by discretizing it into a multiresolution tetrahedral mesh known as a hierarchy of diamonds. This model enables us to extract crack-free approximations of the 3D image at any uniform or variable resolution, thus reducing the size of the data set without reducing the accuracy.

A 3D intensity image is a scalar field (the intensity field) defined at the vertices of a 3D regular grid and thus the graph of the image is a hypersurface in R^4 . We measure the discrete distortion, a generalization of the notion of curvature, of the transformation which maps the tetrahedralized 3D grid onto its graph in R^4 .

We evaluate the use of a hierarchy of diamonds to analyze properties of a 3D image, such as its discrete distortion, directly on lower resolution approximations. Our results indicate that distortion-guided extractions focus the resolution of approximated images on the salient features of the intensity image.

1 Introduction

The increasing sizes of 3D images require multiresolution models that allow the extraction of adaptive representations of the image at many different levels of detail. A very large class of multiresolution models is provided by nested meshes, in which all elements are defined by the uniform subdivision of a small set of cells. Examples include: *octrees* formed by cubes and nested tetrahedral meshes based on the *Longest Edge Bisection (LEB)* operation.

A 3D image I is typically represented as a cubic grid of samples, called *voxels*, encoding the image intensity at each grid point. Since this representation requires the entire domain to be uniformly sampled at the resolution of the smallest feature within the image, *octrees* are often used to aggregate the featureless regions of the image into a nested hierarchy of cubes.

A disadvantage of voxel-based cubic representations

is that cracks between adjacent cubes lead to discontinuities over the image's intensity function. This can be remedied by using a dual representation where the grid points are vertices of a crack-free tetrahedral mesh decomposing the domain of the image.

In our work, we use nested tetrahedral meshes based on the *Longest Edge Bisection (LEB)* operation, in which a tetrahedron t is bisected along the plane defined by the midpoint of its longest edge e and the two vertices of t not adjacent to e . These were originally introduced for domain decomposition in finite element analysis [4], and have since been applied in many different contexts, including scientific computing, surface reconstruction and volume segmentation.

The curvature of a shape is an important property that facilitates the study of its topological and metric properties, as it measures the local defect from flat Euclidean space. There is a rich literature dealing with the problem of defining and computing *discrete* curvature estimators for triangulated surfaces [2]. A particularly simple and efficient method is *concentrated curvature* [1].

Discrete distortion [3] is a recent evolution of the notion of concentrated curvature for three-dimensional tetrahedralized shapes embedded in 4D space. Since a 3D intensity image I is a scalar field associated with the vertices of a cubic 3D grid, the graph of I is a hypersurface in R^4 and a tetrahedral mesh Σ decomposing the 3D cubic grid is embedded in R^4 . We can view the intensity values of I as constraints on the vertices of Σ , which induce a distortion of geometry on the tetrahedral hypersurface representing the graph of I in R^4 . This distortion can be estimated through trigonometric relations linking dihedral and solid angles of the tetrahedralized hypersurface. As for surface curvature, discrete distortion highlights the local curvature of the graph of the intensity image, i.e. of the tetrahedralized hypersurface.

The contribution of this paper is two-fold. First, we demonstrate the advantages of a multiresolution tetrahedral model for the representation of a large 3D intensity

image I . This model enables the extraction of crack-free approximations to I on which we can directly analyze properties of I , such as its distortion. Second, we explore the viability of using distortion as a tool for 3D image analysis. Specifically, we generate a multiresolution model based on the distortion as well as on the intensity value.

2 Diamond Hierarchies

A *Hierarchy of Diamonds* Δ is a multiresolution representation of a 3D image I based on clusters of tetrahedra, called *diamonds*, sharing the same bisection edge, that we call the *spine* of the diamond. We have studied the theory and the properties of hierarchy of diamonds in arbitrary dimensions [5].

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 over the diamonds defines a direct dependency relation over the samples of I , which can be encoded as a *directed acyclic graph* (DAG).

Due to the regularity of the vertex distribution and the subdivision rule, this model generates only three similarity classes of diamonds containing well-shaped tetrahedra [4]. 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 I . From the coordinates of the central vertex, we use bit manipulations to extract the complete parent-child relations. Thus, a hierarchy of diamonds is encoded as the collection of the central vertices of its diamonds.

We can have a *full* or a *partial* hierarchy of diamonds representing a 3D image. A *full hierarchy of diamonds* represents the complete 3D image and thus the coordinates of its central vertices can be implicitly encoded, while a *partial hierarchy of diamonds* requires an encoding of the coordinates of the central vertices. A partial hierarchy can be especially useful when we are interested in only a sparse subset of a large image, or when the image is oversampled.

A (full or partial) hierarchy of diamonds Δ is used to efficiently extract variable-resolution tetrahedral meshes Σ approximating a 3D image I 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 image I . In contrast to the octree-based approaches, such tetrahedral meshes

have a higher degree of adaptability to the selection criterion, and are guaranteed to be free of cracks.

Figure 1a illustrates a variable-resolution tetrahedral mesh Σ containing 77 K vertices and 427 K tetrahedra extracted from a $256^2 \times 161$ sample 3D image representing a Micro CT scan of a tooth. Even though Σ contains fewer than 1% of the samples of the full image, it suffices to extract accurate isosurfaces (see Figure 1b).




Figure 1. (a) Tetrahedral mesh Σ extracted from a 3D Micro CT scan of a tooth. (b) Isosurface embedding within Σ .

3 Discrete Distortion

The graph representation of the intensity function f defined on a 3D image discretized as a tetrahedral mesh Σ is a tetrahedral hypersurface $(\Sigma; f)$ in R^4 , whose distortion with respect to the underlying discretized cubic grid is due to the effects of the intensity values. Similarly to the concentrated curvature of triangulated surfaces in 3D space [1], one may compare the solid angle defect at vertices of tetrahedralized hypersurfaces in R^4 . We define the *distortion* at a vertex p as

$$D(p) = \sum_{t_i \in t(p)} s_i - \sum_{T_i \in T(p)} S_i, \quad (1)$$

where s_i (resp. S_i) is the solid angle within Σ (resp. within $(\Sigma; f)$), at vertex p of tetrahedron t_i (resp. T_i), and $t(p)$ (resp. $T(p)$) is the set of tetrahedra incident to p in Σ (resp. $(\Sigma; f)$). For internal vertices the sum $\sum_{t_i \in t(p)} s_i$ reduces to 4π . Theoretical properties of discrete distortion are studied in [3], where the concept of distortion for the edges of the mesh is also formalized.

It can also be shown that discrete distortion gives positive values to locally convex or concave areas of the tetrahedralized hypersurface; negative values to saddles; and null values to flat areas.

Another relevant property is that distortion is mesh-dependent, i.e., the distortion at a vertex depends on the

way in which its neighborhood is triangulated. Even if the local geometrical shape of Σ does not change in R^3 after a remeshing, its image $(\Sigma; f)$ in R^4 changes and, thus, distortion changes to quantify this effect. The mesh-dependent property highlights the effect of a multi-resolution representation.

As an example of distortion, the electron density of the Hydrogen molecule H_2 (Figure 2, top row) and its associated discrete distortion (Figure 2, bottom row) are represented through three pairwise-orthogonal planar slices of the molecule, colored blue for low values and red for high values. High density values occur within spheres centered at the atom loci, at the (symmetry) center of the molecule (i.e., the equilibrium position) and within nested tori around the symmetry center. We see from the distortion slices that the electron density increases with a uniform speed near the atoms (i.e., almost constant distortion values) and decrease almost linearly (i.e., nearly null distortion values) far from the atoms. The transition of electrons from the spheres to the symmetry center happens abruptly, as can be seen at blue points on the distortion slices between atom’s positions and the center. This information is not captured directly from the electron density image.




Figure 2. Axis aligned planar slices of 3D image representing Hydrogen electron density (top) and distortion (bottom).

4 Experimental Results

We demonstrate the interplay between mesh resolution and discrete distortion through two experiments.

First, we illustrate the role of the mesh resolution in its computed distortion over a series of increasingly fine meshes. For this, we generate a diamond hierarchy Δ_H based on the intensity values of the 128^3

Hydrogen dataset, described in Section 3, which contains 2.1 M vertices. The error of a diamond δ is computed as the maximum difference between the intensity values of all grid points within δ , and the value obtained by linear interpolation over δ 's vertices. We extract a series of meshes Σ_{ϵ_i} of uniform approximation error ϵ_i from Δ_H , using threshold values of $\epsilon_i \in \{30\%, 10\%, 5\%, 2\%, 1\%, 0.5\%, 0\%\}$ of the total error and then evaluate the distortion of the vertices of these meshes. The number of vertices in these meshes are $\{.2K, .5K, .7K, 2K, 7K, 20K, 544K\}$, respectively.

Figure 3 shows the cumulative distribution function (CDF) of the discrete distortion D (horizontal axis) of the vertices of each mesh. The sharp spike in the CDF of all datasets around distortion $D(v) = 0$ 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. For example, when $\epsilon = 0$, more than 99% of the vertices have absolute-valued distortion $D(v) \leq |.2|$. This indicates 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.






Figure 3. Cumulative distribution functions of Hydrogen distortion.


In our second experiment, we use distortion to guide the extraction of lower resolution approximations to the full image. Here, we consider a 32^3 image representing the electron density of the C_{60} Buckyball molecule. As for the Hydrogen molecule, the density corresponds to the locus of atoms around which electrons revolve, as well as the regions that they avoid. For this experiment, we compute the distortion of the mesh Σ_B at full resolution (containing 36 K vertices and 216 K tetrahedra). Since the vertices of Σ_B correspond to a regularly sampled scalar field, we generate a diamond hierarchy Δ_B based on the distortion values rather than the inten-



(a) 50% distortion error ($|V| = 3k$, $|T| = 15k$)



(b) 5% distortion error ($|V| = 17k$, $|T| = 90k$)



(c) 0% distortion error ($|V| = 34k$, $|T| = 189k$)

Figure 4. Slices of tetrahedral meshes extracted from the Buckyball dataset colored according to distortion (left) and intensity (right) for error thresholds of 50% (a), 5% (b) and 0% (c).

sity values. We then extract several variable resolution meshes from Δ_B guided by Σ_B 's distortion. We visually compare several extracted datasets of relative error $\epsilon = \{50\%, 5\%, 0\%\}$ in Figure 4 through axis-aligned slices colored by the distortion (left), where positive distortion is colored red and negative distortion is colored blue. We also visualize the intensity values of the image over the same meshes in Figure 4 (right).

By comparing the left and right parts of Figure 4a, we observe that the resolution is concentrated initially on the regions with higher absolute-valued distortion, such as the dark red regions corresponding to the carbon atoms of the molecule.

Since vertex distortions correspond to singularities in the graph of the image, there is a relatively high interpolation error for the regions containing such singularities. Thus, the salient features of the intensity field are well approximated, even at extremely coarse approximations, while further increases to the resolution enhance the approximation of the intensity function. We have obtained similar results on other 3D images describing medical and scientific datasets.

5 Concluding Remarks

We have demonstrated that a multiresolution model based on clusters of tetrahedra, called diamonds, enables the analysis of a 3D image through crack-free approximations encoded as tetrahedral meshes. We have also demonstrated the utility of discrete distortion in analyzing the approximated images and in guiding the extraction process to yield accurate approximations of the original image.

One of the important aspects of mesh-based multiresolution models is that we can analyze the image using much fewer samples than the full image. This facilitates our analysis of large images using significantly fewer resources.

References

- [1] P. Aleksandrov. *Topologia Combinatoria*. Torino, 1957.
- [2] T. Gatzke and C. Grimm. Estimating curvature on triangular meshes. *Int'l J. on Shape Modeling*, 12:1–29, 2006.
- [3] M. M. Mesmoudi, L. De Floriani, and U. Port. Discrete distortion in triangulated 3-manifolds. *Computer Graphics Forum*, 27(5):1333–1340, 2008.
- [4] M. C. Rivara and C. Levin. A 3D refinement algorithm for adaptive and multigrid techniques. *Communications in Applied Numerical Methods*, 8:281–290, 1992.
- [5] K. Weiss and L. De Floriani. Diamond hierarchies of arbitrary dimension. *Computer Graphics Forum (Proceedings SGP 2009)*, 28(5):1289–1300, 2009.