C^1 continuities detection in triangular meshes

Luca Di Angelo, Paolo Di Stefano *

Università degli Studi de L'Aquila, Italy

ARTICLE INFO

ABSTRACT

The identification of C^1 continuities is important in many applications involving point clouds or triangular meshes, such as surface segmentation, inspection and rendering. The methods in literature have some limitations which make them strongly dependent on some properties of the mesh (point typology, mesh resolution, uniformity of the shape of triangles and error in point location). Furthermore, some of them do not discriminate non-regular points from those that are inside a band around them.

In this work, a new method for automatic detection of C^1 continuities in triangular meshes is presented. The method introduces an original function, called *sharpness indicator*, which enables us to evaluate properties related to surface smoothness. The performance of the new method is compared with that of four methods presented in literature as regards the recognition of C^1 continuities both in synthetic and real meshes. Results are analysed and critically discussed.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

In a continuous surface, a C^1 discontinuity pertains to a point whose neighbourhood cannot be represented by a smooth function. Typically, these points are those pertaining to the edges of the geometrical model where two different non-tangent surfaces meet; in literature they are also referred to as *sharp* or *ridge points*. C^1 discontinuity can also be detected in isolated points associated with surface defect, which are also defined as *corner points*.

In the past few years the concept of *sharp point* has been extended to discrete geometrical models due to the more widespread use of a representation scheme based on triangular meshes or point clouds. In these cases the main problem lies in an intrinsic C^1 discontinuity of the points pertaining to a non planar surface of a tessellated geometric model. At present, the identification of C^1 discontinuities is an important and preliminary stage in many different applications involving geometric models defined by point clouds or triangular meshes, such as:

- *Point cloud segmentation*, where *ridge point* detection makes it possible to determine the boundary edges between the surfaces in the *edge-based* algorithms [1–3] and to constrain the growing process in the *hybrid* segmentation approaches [4–6];
- *Tolerance inspection*, where the identification of *corner points* is necessary in order to identify and eliminate singularities having a large noise content which makes verification unreliable; furthermore, the detection of the *sharp points* enables us to simplify and speed up the registration process, where a scanned object is aligned with its CAD model [7,8];

- *3D geometric model rendering*, in the point-based rendering techniques; in computer graphics, the recognition of C^1 discontinuities is essential to obtain a robust and reliable rendering [9]; furthermore, *ridge point* detection is necessary to mark the edges of the model.

The concept of C^1 surface discontinuity is associated with the mathematical concept of differentiability of the function that defines the geometrical surface. In the continuous case, since the analytical formulation of the surfaces is known, C^1 discontinuity is clearly recognisable. That is the case of geometric models defined by NURBS or other kinds of parametric surfaces. In those geometric models defined by point clouds or by tessellated surfaces, C^1 continuity has been lost due to the discretisation process and it cannot be assumed to be a property of the surface but rather must be recognised by analysing the data retained in the location of the points.

The concept of surface C^1 continuity recognition in tessellated models is crucial since many of the methods presented in literature identify C^1 discontinuities by evaluating some differential geometric properties. The differential geometric properties can be evaluated only when the *surface regularity* [10] at each point has been preventively recognised. Therefore, due to the nature of a tessellated model, it is rather more appropriate to recognise surface continuity than C^1 discontinuities, which are an intrinsic characteristic of any non-flat tessellated surface.

In literature, some algorithms are presented to identify C^1 discontinuities starting from a mesh [7,11–23]. These algorithms, however, do not satisfactorily solve the problem in practical situations. The recognition process is blurred by the polyhedral conformation of the geometric model, by noise due to errors in the process of acquisition of the model and by the algorithms adopted

for surface tessellation. Some algorithms have limitations due to the criteria that must be adopted to choose the threshold value that discriminates discontinuities. It depends on many factors such as point typology (umbilical, elliptical, hyperbolic, parabolic), mesh resolution and uniformity, and noise coming from the point acquisition process. Other algorithms do not exactly distinguish the *sharp points* from those which are inside a band around them.

In order to solve these problems, in this work a new method is proposed. It focuses on the recognition of C^1 continuities through the evaluation of a specifically defined function (*smooth function*). This method is tested in both synthetic and practical tessellated surfaces. The test cases are the typical benchmarks used in the related literature. The results derived from these experiments are critically discussed hereinafter.

2. Related works

It is on account of the importance of the problem that numerous algorithms to identify C^1 discontinuities are presented in literature. They analyse some geometrical properties recognised in the neighbourhood of the point under examination. The most common algorithms can be classified based on the type of geometrical parameters analysed and used to define the indicator suited to recognise discontinuities:

- Error in local least-squares fitting by a plane or by a quadric surface;
- Normal vectors;
- Curvatures processed in different ways.

The above-mentioned geometrical parameters are evaluated at each data point (or in the triangles in the neighbourhood of the point) and properly elaborated; they furnish the values of the indicators of surface discontinuities. C^1 discontinuities are detected as those having an indicator value over a properly defined threshold.

Benko et al. in [11] present, as indicator, a normalised value of the error in the approximation of the neighbourhood of each data point with a plane (*planarity*). Any point is defined as sharp if the *planarity* value is higher than a threshold value. This method is founded on the concept that a smooth surface can be locally approximated by the tangent plane. Since a tessellated surface is a discrete representation of a curved surface, the functionality of this method is affected by two factors: the local mesh dimension and curvature. A high value of mesh dimension or curvature gives rise to a high value of the normalised error. The normal vector based methods [7,12–17] identify C^1 discontinuities by comparing the dot product of the normals of two triangles sharing an edge or their dihedral angle with a reference value. Moreover, these methods do not take into account mesh sizes and curvatures. With a view to identifying sharp points in noised meshes, Wang in [17] proposes a method which improves normal vector deviation methods by preprocessing the original mesh by several passes of bilateral filtering. The strong point of this method is precisely the pre-processing filtering and not so much the robustness of the indicator used to detect the discontinuities, which shows the typical problems of the normal vector deviation indicators.

In order to overcome the limitations of normal vector based methods, Wang in [18] proposes a different approach and introduces a new detector which measures the smoothness of a tessellated surface, defined as follows:

$$\tau\left(\mathbf{p}\right) = \frac{1}{\lambda}\max\left(W\left(\mathbf{p}, T_i\right) \cdot W\left(\mathbf{p}, T_j\right) \cdot \left(1 - \boldsymbol{v}_i \cdot \boldsymbol{v}_j\right)\right) \qquad (1)$$

where:

- $\lambda = \begin{cases} \mu_{mesh} & \text{for regular meshes} \\ 1.5 \cdot \mu_{mesh} & \text{for irregular meshes;} \end{cases}$
- μ_{mesh} is the mean size of the mesh edges;
- $W(\mathbf{p}, T_i)$ is the window function $= \begin{cases} 1 & d(\mathbf{p}, T_i) \leq \lambda \\ 0 & d(\mathbf{p}, T_i) > \lambda; \end{cases}$

- $d(\mathbf{p}, T_i)$ is the Euclidean distance between $\mathbf{p}$ and T_i;
- $\boldsymbol{v}_i \cdot \boldsymbol{v}_j$ is the dot product of the unit normal vectors of the triangles T_i and T_j.

The point is defined as sharp if $\tau(\mathbf{p})$ is higher than a threshold value (κ). This method takes into account the mesh dimension by making use of the factor λ that weights in $\tau(\mathbf{p})$ the dot product of the normals of the triangles selected by the window function.

Jiao and Alexander in [19] use the normal vectors of the triangles incident on the point under analysis in order to build the following matrix:

$$\mathbf{A} = \sum_i w_i \cdot \boldsymbol{v}_i \cdot \boldsymbol{v}_i^T \qquad (2)$$

where w_i and $\boldsymbol{v}_i$ are, respectively, the area and the normal of the *i*th triangle belonging to the 1-ring neighbourhood. The authors characterise the local smoothness of the point by means of the relative sizes of the eigenvalues of the matrix $\mathbf{A}$ (λ_1, λ_2 and λ_3, with $\lambda_1 \geq \lambda_2 \geq \lambda_3$). In particular, this method distinguishes the points as follows:

- corner, if $\lambda_3 \geq \beta \cdot max(\lambda_1 - \lambda_2, \lambda_2 - \lambda_3)$ or $\mathbf{A}$ has full rank;
- ridge, if $\lambda_2 \geq \alpha \cdot \lambda_1$ or $\mathbf{A}$ has rank two;
- smooth, if other or $\mathbf{A}$ has rank one.

α and β are two reference values which are evaluated from the threshold value Ψ and defined as follows: $\beta = \cot(\Psi)$ and $\alpha = \tan^2(\Psi/2)$.

In order to develop a stable procedure to determine edges for which there exist discontinuities, both C^1 and C^2, Jiao and Bayyama in [20] present a complex method specially designed for synthetic meshes. It starts with the labelling of the edges and vertices of the mesh by analysing four indicators (angle defect, turning angle, one-sided turning angle and dihedral angle), all of which are based on the elaboration of the normal at vertices and triangles; the labelling is carried out by comparing the indicators with eight different thresholds whose values are experimentally determined. With a view to removing the numerous false positives recognised by the method, the authors develop further algorithms. These require a further edge-labelling process, based on the above-mentioned indicators, which carries out a classification and filtration of "candidate curves" to obtain final sequences of edges. This filtration process makes it difficult to apply the method to meshes coming from scanned objects; in these cases, and because of both the acquisition errors and the objects' singularities, points are not perfectly located on a ridge or a tangent line.

Other researchers identify C^1 discontinuities through curvature analysis. Baker [21] detects C^1 discontinuities by comparing weighted principal curvatures with a properly chosen threshold value. The weight factor takes into account the mean size of the 1-ring neighbourhood of the point under examination. Vieira and Shimada [22], by using an empirical formula, compare the largest estimated principal curvature (k_1) at the point examined with the average length of the edges incident on it ($\mu_{mesh,p}$). A point $\mathbf{p}$ is identified as sharp when:

$$\mu_{mesh,\mathbf{p}} \cdot |k_1| > threshold\ value. \qquad (3)$$

The authors in [22] assume the *threshold value* to be equal to 0.1.

These methods [21,22] only process data evaluated at the mesh points, so they cannot exactly differentiate the nodes belonging to a sharp edge from those that are inside a band around it. This limitation is overcome by using algorithms which analyse the mesh edges. Huang and Menq in [3], compare the directional curvatures evaluated at both vertices of each edge of the mesh. When the absolute difference of these curvatures is over an assigned threshold value, the edge is assumed to be sharp and so are its related vertices. Since this threshold value depends also on the typology of the neighbourhood, its choice is often difficult. Finally, as pointed