ELSEVIER

Topological recognition of polyhedral objects from multiple views

Aldo Laurentini

Dipartimento di Automatica ed Informatica, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

Abstract

The aspect graph (AG) has been proposed as a viewer-centered tool for object recognition. Although the aspects could be used for recognition from one image, the full information stored in the AG (aspects and visual events) suggests a *multiple* views, possibly active, approach. The topological nature of the AG also suggests a *topological* match of images of the unknown object and stored aspects.

In this paper we present a theoretical investigation on the use of the AG for topological recognition from multiple views of polyhedral objects. First, we discuss the topological matching process, and give a suitable topological definition of aspect. Moreover, since topological identification is approximate, we tackle the problem of understanding the ability of the AG to topologically discriminate different objects. More precisely, we address the question: how "similar" are two polyhedra with the same AG? The isomorphism of polyhedra is chosen as the reference similarity condition. The cases of general and convex polyhedra are discussed under perspective and parallel projections. In practice, it could be difficult to identify the visual events from the images of the unknown object. Thus, we also compare isomorphism and the similarity induced by: (i) a topologically reduced aspect graph; (ii) the set of topologically different aspects, neglecting visual events. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Aspect graphs; 3-D object recognition; Computer vision; Polyhedral objects; Topology; Isomorphism

1. Introduction

The *aspect graph*, introduced by by Koenderink and van Doorn [19] under the name of *visual potential*, has been the object of intensive study in the last years.

E-mail address: laurentini@polito.it (A. Laurentini).

The idea at the basis of the AG is to cluster the infinite possible views of an object into a finite set of representative views, or *aspects*. In principle, an aspect could be any collection of features of a 2-D image, invariant or relatively invariant under some limited displacement of the viewpoint. Aspects defined according to this general approach have been used by Seibert and Waxman [29]. However, in most papers on this subject the views to be clustered are line drawings generated by projecting onto the image plane the edges of polyhedral objects, or the "creases" (surface normal discontinuities) and "limbs" (depth discontinuities) of general objects.

Two approaches have been used for selecting the representative views. According to one approach (see for instance [10,15,20]), the viewing space is sampled by equally spaced viewpoints. The second approach consists in clustering the line drawings according to their *topological structure*. The topological clustering, originally suggested by Koenderink and van Doorn, leads to a partition of the viewing space into non-uniform viewing regions, each corresponding to a particular aspect, determined by the geometry of the object.

The *topological adjacency* between aspects, consequence of the topological clustering, allows to arrange the aspects into a graph structure, the AG precisely. In the AG each node is labeled with an aspect and each arc represents a topological change, referred to as a *visual event* or *transition*. Nodes and arcs are associated to the viewing regions and their boundaries respectively.

Aspects and visual events could be related to the biological vision process [29].

Many algorithms have been developed for computing the topological partition of the viewing space for various categories of objects under parallel or perspective projection. More specifically, polygons, convex and general polyhedra, articulated assemblies and curved surface objects have been considered. For further detail, the reader is referred to the survey paper [3] and to [4,7,8,12,21,22,27,28,30].

The AG approach to object identification is not without shortcomings. A discussion between critics and supporters of the AG idea can be found in [9]. One problem is the huge amount of aspects of even relatively simple objects. To cope with it, the idea of finite resolution AG has been proposed [7,30]. According to another approach, the recognition of a complex object is performed by identifying in its line drawings sub-images of simple 3-D primitives [6].

Anyway, a certain unbalance can be remarked between the efforts aimed at computing the aspect graphs and the practical applications reported. The author feels that the full power of the aspect graph idea has not been fully exploited.

First, the full information stored in the AG, that is aspects and transitions, suggests a *multiple views approach*, possibly in relation with an active inspection algorithm. However, little work in this area has been reported. Seibert and Waxman [29] presented a recognition algorithm along this line, though their definition of aspect is not related to line drawings. Some pioneering work based on multiple line drawings of polyhedra was made by Underwood and Coates [32], before the introduction of the very concept of aspect graph.

Second, *topologically matching* images of an unknown object and stored aspects appears a most straightforward consequence of the topological nature of the aspect graph. However, even if a few recognition algorithms make use of this idea (see for instance [6]), the capabilities of the topological approach have not been thoroughly investigated.

The purpose of this paper is to theoretically explore these two ideas. More precisely, we will study *the recognition of polyhedral objects by topologically matching multiple views to stored aspects and transitions.*

First we address the problem of topologically matching aspects. After discussing the available topological definitions of aspects, we propose a definition convenient for any possible aspect of a general polyhedron, and show that the corresponding topological matching is in most cases a simple process. This is a point in favor of topological recognition, in view of the impressive growth of the number of aspects of a polyhedron with the number n of faces.

Of course, a topological match is something less than a match based on metric properties. Thus we also address the problem of understanding what can be inferred about the shape of a topologically identified object. In other words, we tackle the problem of finding *how similar are two objects with topologically equivalent aspect graphs*. For doing this, we choose the *isomorphism* of polyhedra as reference similarity condition. We compare with isomorphism the similarity induced by having the same AG under various conditions (parallel or perspective projections, convex or general polyhedra). It results that having the same AG implies in some cases a similarity stronger than isomorphism.

In practice, when analyzing multiple images of an objects it could be difficult:
 (i) to identify the corresponding vertices and edges of line drawings extracted from different frames;
 (ii) to understand whether two line drawings are actually adjacent.

This suggests to reduce the information stored by merging together all the topologically equivalent aspects. Thus, we also compare isomorphism and the similarity induced by:
 (i) a topologically reduced AG;
 (ii) the set of all the topologically different aspects, neglecting transitions.

This paper is organized as follows. In Section 2 we summarize the relevant matter on the AG of polyhedra. In Section 3, we discuss the topological definition of aspects and their comparison. In Section 4 we introduce isomorphism as the reference condition for topological similarity. The discriminating power of the perspective AG for general polyhedra is discussed in Section 5. The case of the parallel AG is considered in Section 6. The similarity of polyhedra with the same reduced AG or aspect set is analyzed in Section 7. The results obtained are summarized and discussed in Section 8.

2. The aspect graph of polyhedral objects

The graph structure of the aspect graph is the dual graph of the partition of the viewing space. For a convex polyhedron S, the edges are completely visible or invisible. The topological structure of the line drawing changes when the viewpoint crosses each plane supporting a face. These n planes divide $\mathbb{R}^3 - S$ into bounded and unbounded viewing zones. The *perspective* aspect graph has a node for each zone and an arc for each face of the partition. The *parallel* aspect graph (always planar) is only relative to the unbounded zones. A popular equivalent formulation makes reference to the partition of the gaussian sphere. n planes divide $\mathbb{R}^3 - S$ into $O(n^3)$ and the gaussian sphere into $O(n^2)$ viewing

regions. These are the bounds for the numbers of aspects and transitions of the perspective and parallel aspect graphs of convex polyhedra.

For non-convex polyhedra, the edges could be visible only partially. The topological structure of the line drawing also changes when the viewpoint crosses VE- and EEE-surfaces, originated by self-occlusion. The VE-surfaces are generated by the lines passing through a vertex and an edge. They account for topological changes due to an edge which occludes a vertex or a vertex which occludes an edge, and are $O(n^2)$. The EEE-surfaces are due to an edge which occludes a T-junction or vice-versa (a T-junction is a point where the projections of two non-adjacent edges meet). EEE-surfaces are formed by the lines passing through three edges, and thus are $O(n^3)$. All these surfaces divide the gaussian sphere into $O(n^6)$ zones, and $\mathbb{R}^3 - S$ into $O(n^9)$ zones. Plantinga and Dyer [28] presented examples of worst case non-convex polyhedral scenes.

3. Topologically matching aspects

With the expression *topological recognition from multiple views* of a polyhedron we refer to the following approach. We are given a set of aspect graphs, one for each object which could be observed. From a 2-D image of the object we derive a line drawing that we topologically compare with the stored aspects (the topological comparison will be exactly defined in the following). If one or more matches are found in one AG only, the object has been recognized. Otherwise, we continue observing the object, supposed in relative motion with respect to the viewpoint. This relative motion could be driven by an active recognition algorithm. From the changing 2-D image we derive a sequence of topologically distinct line drawings which we attempt to match with sequences of adjacent aspects in the various AGs, starting from the first match found. As soon as a match becomes possible for only one AG, the object is identified. Since the topological identification is approximate in nature, it could be the first step of the whole recognition process.

In the rest of this section we will focus on the topological definition of aspect and the topological matching process.

3.1. The topological definition of aspect

The idea is to match two aspects if they are "topologically equivalent". What does exactly this means, or, in other words, which topological definition of aspects must we assume? Several topological definitions of aspect have been proposed, but we will show that most of them are *ambiguous*, i.e., they can be equal for line drawings which do not conform to the intuitive idea of topological equivalence. Anyway, the available definitions do not cover all the possible cases of polyhedral line drawings.

An aspect is most often defined as the *graph structure*, labeled in some way, of the *image structure graph* (ISG), i.e., the planar graph which has a node for each junction and an edge for each segment of the corresponding line drawing. The ISG is invariant for each viewing zone, but one simple example shows that it is not adequate to our purposes. Let us consider the line drawings LD and LD′ in Fig. 1.
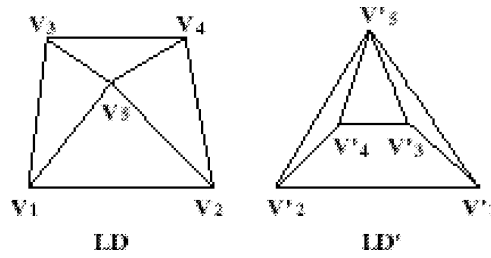
Fig. 1. The line drawings LD and LD′ have isomorphic ISGs, but do not satisfy the idea of topological equivalence.
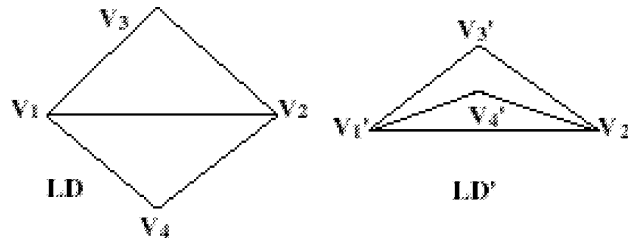

Fig. 2. Another example of line drawings having isomorphic ISGs which do not appear topologically equivalent.

Both LD and LD′ can be interpreted as images of convex polyhedra. Their ISGs are isomorphic (the corresponding vertices are shown in the figure), but it is clear that LD and LD′ cannot be considered instances of topologically equivalent aspects. A way for finding other cases of ambiguous ISG descriptions is to consider two *Schlegel diagrams* of a polyhedron P. A Schlegel diagram of a (transparent) convex polyhedron P relative to a face f is a line drawing obtained by projecting onto f, from a suitable viewpoint, all the edges of P [13, p. 42]. It is clear that all the Schlegel diagrams of a polyhedron, which can always be interpreted as line drawings of opaque convex polyhedra, have isomorphic ISGs. The line drawings of Fig. 1 are Schlegel diagrams of a square pyramid, relative to the base face and a side face.

If we drop the convexity restriction, which requires the non-perimetrical vertices of the line drawing to have degree three at least, we can find even simpler examples of ambiguous ISG descriptions, like that of Fig. 2.

As a matter of fact, the ISG representation is not even sufficient for distinguishing adjacent viewing zones, at least for origami-like objects. In fact, if we assume that the triangular face specified by vertices V_1, V_2 and V_4 in Fig. 2 has zero thickness, LD and LD′ can represent adjacent aspects of the same object.

According to a definition proposed by Gigus and Malik [12], we could label the edges of the ISG representing concave, convex and occluding edges by the Huffman–Clowes scheme. Actually, this labeling can resolve some ambiguities, for instance those of Fig. 1 and Fig. 2, but not any ambiguity, as shown by a more complex example (Fig. 3).
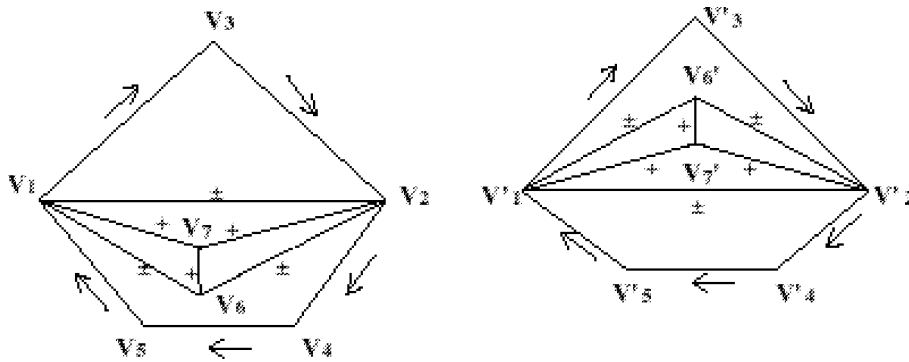
Fig. 3. Two line drawings with isomorphic ISGs which cannot be distinguished by the Huffman–Cloves labeling scheme.

In addition, comparing labeled aspects would require to label the line drawings of the unknown object, a hard problem which in general is NP-complete, and can produce many different solutions (see [18]).

Another proposed labeling scheme, relative to the type of junctions (see [3]), can be easily shown to be ambiguous. In any way, it is not in keeping with our purely topological approach, since it is based on the value of the angles between the edges.

Rewording our problem according to the graph theory, the line drawings considered so far are rectilinear planar embedding of bi-connected planar graphs. A planar embedding is a physical representation of the graph structure such that no two edges cross each other. The intrinsic ambiguity of the ISG representation is related to the fact that bi-connected planar graphs can have different planar embeddings. This statement also holds for rectilinear embeddings, since any planar embedding with curved edges can be transformed into a rectilinear embedding [1]. In the literature a planar embedding is specified with the *ordered adjacency lists* of its vertices (see [5,17]). It is immediate to verify that this description of a line drawing is stable within the viewing regions and changes only crossing their boundaries. It is also easily seen that the ordered adjacency lists are different for the pairs of line drawings in Fig. 2 and Fig. 3. Unfortunately however, as it is even this topological description could be ambiguous. In fact, let us consider again the two line drawings LD and LD$'$ of Fig. 1. It is easy to verify that their descriptions in terms of adjacency lists are equal.

In the following we introduce a topological definition convenient for the topological recognition and covering any possible case of polyhedral line drawing. Let us consider for the moment connected line drawings, as those in Fig. 1, 2 and 3. Any such line drawing, relative to convex, concave or even origami-like objects, can be seen as a set of adjacent polygons, each representing the projection of a face or part of it. Thus, a natural topological description, that we call TD, consists in specifying an aspect as a set of polygons, and each polygon as the list of its vertices in clockwise order. This simple idea is the same that leads to the face graph (see [6]). TD satisfies the intuitive idea of topologically equivalent line drawings, and solves ambiguity problems like those reported. For instance, LD of Fig. 1 is specified as $((V_3, V_4, V_5), (V_3, V_5, V_1), (V_1, V_5, V_2), (V_5, V_4, V_2))$, while LD$'$ is specified

$A=((V1, V2, V3, V4), (V4, V3, V5, V6), (V1, V4, A', V6, V7))$
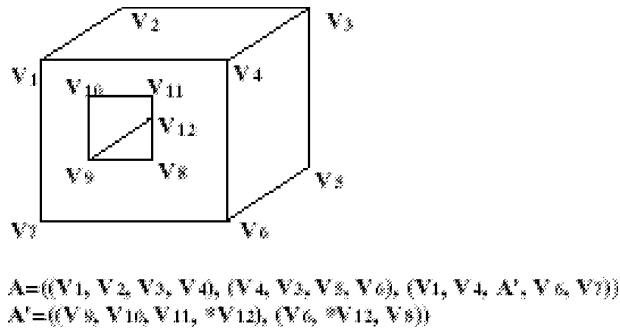$A'=((V9, V10, V11, *V12), (V6, *V12, V8))$

Fig. 4. An aspect made of two not connected line drawings, and its TD.

as $((V'_1, V'_2, V'_4, V'_3), (V'_1, V'_3, V'_5), (V'_5, V'_3, V'_4), (V'_5, V'_4, V'_2))$. The distinctive features of TD are the following (we omit formal proofs of these rather obvious statements):

– TD is stable within each viewing region and changes when the viewpoint crosses the boundaries of the region.
– From TD we can *uniquely* derive (in linear time):
  (1) the ISG description;
  (2) the ordered adjacency lists for each node;
  (3) the list, in clockwise order, of the vertices composing the external polygon.

Observe that the information under points (2) and (3) is equivalent to TD. Point (2) is particularly important, since our main theorem (Proposition 2 in Section 5) requires a topological definition of aspect which specifies the ordered adjacency lists.

In the rest of this subsection, TD will be extended to take into account any possible aspect of a general polyhedron. Some possible alternatives for TD will be also discussed.

Non-connected line drawings

Even a connected object could produce non-connected line drawings, as in Fig. 4. In this case, we specify separately TD for the two non-connected components, and supply the additional information that the smaller component lies inside the polygon specified by the vertices V_1, V_4, V_6, V_7.

T-junctions

In general, some nodes of the line drawing correspond to vertices of the originating polyhedron (let us call them *true* nodes), others (as V_{12} in Fig. 4) are due to a T-junction (*false* nodes). We add to the topological description of an aspect a mark (*) for telling false nodes from true nodes (see Fig. 4). We stress that this is in agreement with the usual assumption of being able to exactly extract line drawings from images of polyhedra. Both extracting line drawings and identifying T-junctions require the capability of identifying precisely the projections of the edges, otherwise we would be unable to identify nodes with degree two.
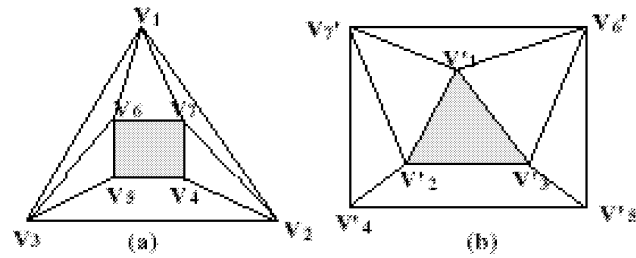
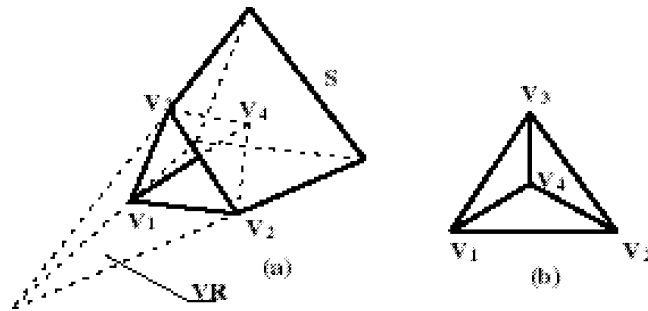Fig. 5. Two line drawings where an internal face is the projection of the background.

Fig. 6. The viewing region **VR** lies in part outside the convex hull of *S* (a); the corresponding aspect (b).

Polygons representing the projection of the background

Unlike the external polygonal boundary, all the internal polygons must be specified in TD, otherwise we produce an ambiguous description. Consider for instance the line drawings in Fig. 5.

Omitting the internal polygon yields equal topological descriptions. Should we distinguish with a particular mark these internal polygons? Obviously, extracting a line drawing from an image implies to identify the external contour of the object; locating internal projections of the background could be much more difficult. In addition, distinguishing these polygons is not necessary for proving Proposition 2. Thus, distinguishing these polygons is only an option of our TD, which could reduce the number of steps required for recognition.

Viewpoints inside the convex hull **CH** *of the polyhedron*

It is clear that we cannot obtain directly a planar complete line drawing from a viewpoint inside **CH** (in practice this is also the case for points near to **CH**). A complete line drawing containing all visible edges can be obtained by projecting the edges onto a suitably small spherical surface centered about the viewpoint. Observe that a viewing region obtained according to the standard definition (see Section 2) could lie in part inside and in part outside **CH**, as shown by the example of Fig. 6.

The corresponding aspect should represent both a planar and a spherical line drawings. An alternative approach consists in distinguishing these two cases by computing the viewing regions using also the surface of **CH** as a boundary. This could reduce the steps of the recognition process without affecting the complexity of the AG (the number of faces
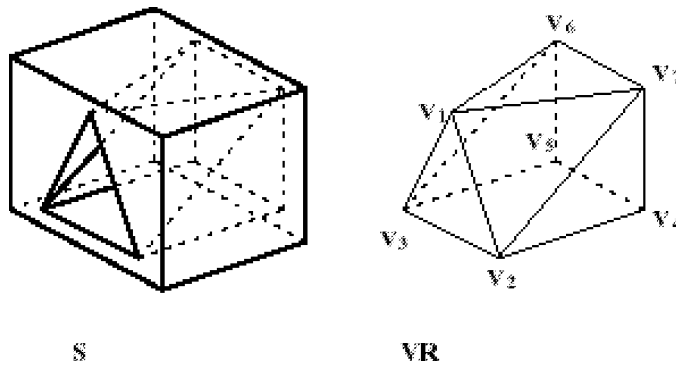
Fig. 7. A non-convex polyhedron *S* and the internal viewing region **VR**.

of **CH** is O(n)). We chose this approach, even if this is not strictly required for the proof of Proposition 2. For viewing regions inside **CH** we apply TD directly to the adjacent "polygons", projections of the faces of the object, of the spherical line drawing, and mark this description as an "internal" aspect.

Some differences with the planar case can arise because of the polygons which are projections of the background. If there is only one such polygon, as in the case in Fig. 6, we can uniquely transform the spherical line drawing into a planar line drawing with the same topological description, where this polygon becomes the external infinite face. If there are several such polygons, the transformation into a planar line drawing with the same topological description is not unique, since each of these polygons can be transformed into the outer face. For instance, the spherical line drawing obtained from the viewing region **VR** inside the object *S* of Fig. 7, cut through by a passing hole, can be transformed into the *two* planar line drawings in Fig. 5 with the same topological description.

Furthermore, observe that:
 (i) the "polygons" which are projections of the background could be small and difficult to distinguish;
(ii) there could be no such polygon at all.

All this suggests a "minimal" TD where all the "polygons" of the spherical line drawing are specified in the same way, and a mark for distinguishing background polygons is only optional. For instance, the aspect of the viewing region **VR** in Fig. 7 is specified by 8 polygons: $((V_1, V_3, V_2), (V_6, V_7, V_4, V_5), (V_5, V_4, V_2, V_3), (V_1, V_6, V_3), (V_1, V_7, V_6),$ $(V_1, V_2, V_7), (V_2, V_4, V_7), (V_3, V_6, V_5))$. This affects the complexity of the topological comparison of aspects, as it will be shown in the next subsection.

Concluding, we have extended the topological definition TD to also take into account:
  (i) unconnected line drawings;
 (ii) T-junctions;
(iii) polygons representing viewing regions inside the convex hull.

Let us complete the topological description of the AG observing that the *transitions* between aspects are implicitly specified by labeling the nodes of each aspect with the same labels of the corresponding nodes of the adjacent aspects. For reducing the number of

topological comparisons, for each aspect we also store: the overall number of edges, true and false vertices, unconnected components, and the number of vertices of the external polygon (if it exists).

3.2. Matching aspects

Let us discuss now the topological comparison between the line drawings obtained from images of the unknown object and the stored aspects. We assume that for both of them the TD description is available. The general idea is to perform the comparison as follows:

(1) Check the overall numbers of edges, true and false vertices, unconnected components and vertices of the external polygon of the two aspects. Stop if they do not agree, otherwise:

(2) Check for isomorphism the ISGs of the line drawing and the stored aspect (if the line drawing is not connected, check the various components, starting from the external, and their relative placement).

(3) If the two graphs have been found to be isomorphic, verify whether the complete topological descriptions TD also match.

A substantial percentage of comparisons should not pass step (1), although to quantify this percentage appears a difficult problem.

In general, the problem of graph isomorphism (step (2)) is not known to be either in the P or NP-complete class (see [26, p. 369]), but efficient algorithms for the planar case have been found. The algorithm proposed by Hopcroft and Wong [14] takes $O(V)$ time (the complexity of graphs algorithms in general depends on both the numbers of edges E and of vertices V, but for planar graphs it is $O(E) = O(V)$). Fast parallel algorithms can be found in [16] and [11]. This is of great importance for the practical feasibility of topological recognition systems, in view of the already mentioned huge amount of aspects for even relatively simple objects.

Step (3) must be performed only after a match in step (2), and thus in a very small number of cases. For determining its complexity, observe first that graphs isomorphism algorithms supply only one isomorphism, but many isomorphisms could exist at the same time between the two ISGs, or, which is the same, each ISG could have many automorphisms.

Actually, their number could depend on V exponentially, as shown by the example in Fig. 8, where the isomorphisms are $2 \times (V - 2)!$, since the $V - 2$ vertices $V_3, \ldots, V_V$ can be associated in all possible ways to $V'_3, \ldots, V'_V$.
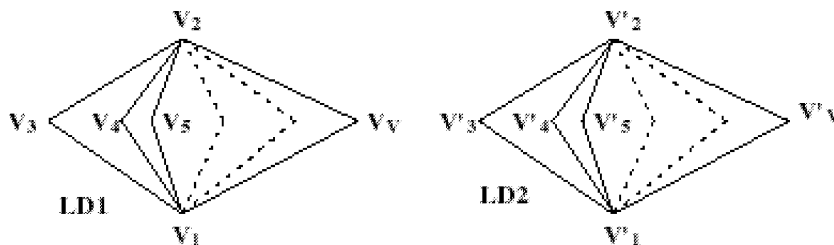

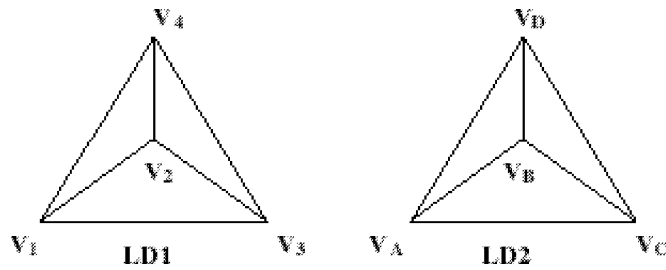Fig. 8. $2 \times (V - 2)!$ different isomorphisms exist between LD1 and LD2, but only one satisfies TD.

Fig. 9. 24 different isomorphisms exist between LD1 and LD2. Three of them also satisfy TD.

Verifying whether the TDs match for one isomorphism takes linear time, but this must be done for each isomorphism and thus a naive implementation of step (3) is exponential in the worst case.

Observe that, in the case of Fig. 8, only two isomorphisms satisfies TD, but this is not always the case. In the example of Fig. 9 three isomorphisms among 24 satisfy TD.

They are:

(1) $V_2 \leftrightarrow V_B$, $V_1 \leftrightarrow V_A$, $V_4 \leftrightarrow V_D$, $V_3 \leftrightarrow V_C$;
(2) $V_2 \leftrightarrow V_B$, $V_1 \leftrightarrow V_D$, $V_4 \leftrightarrow V_C$, $V_3 \leftrightarrow V_A$;
(3) $V_2 \leftrightarrow V_B$, $V_1 \leftrightarrow V_C$, $V_4 \leftrightarrow V_A$, $V_3 \leftrightarrow V_D$.

In cases like this, if the object has not been already identified since we have matched the observed sequence of line drawings with multiple sequences of adjacent aspects in several AGs of the stored objects, we must continue the topological recognition process for all possible cases of correspondences between vertices which satisfy TD. Matching the vertices of adjacent aspects is necessary for guaranteeing that the full capability of discriminating objects of the AG is used. A simpler approach consists in finding only if one topological match exist, without attempting to match the corresponding vertices of two adjacent aspects in the AG with the corresponding vertices of two adjacent aspects of the unknown object. This idea leads to the *reduced* AG, fully discussed in Section 7. We shall see that the reduced AG is less powerful then the regular AG for distinguishing different objects.

Thus, we have seen that the a naive implementation of step (3) of the algorithm is exponential in the worst case. In the following we outline an algorithm which is $O(V^2)$ for the case where an external polygon can be specified.

Let m ($\leqslant V$) be the number of vertices of both outer polygons. There are m possible ways for pairing these m vertices. For each of these m cases, we attempt to match the remaining vertices by using the ordered adjacency lists, starting from the vertices already paired. If the two TDs are equivalent, at the end all the vertices are correctly paired, otherwise some inconsistent pairing occurs for some adjacency list.

An example will show how the algorithm works.

Let us consider the two aspects represented by the line drawings in Fig. 10 (the nodes are all true nodes). Their ISG are isomorphic, their external polygons have 5 vertices so that we must check $m = 5$ cases. Let us consider the case where $V_1 \leftrightarrow V_a$, $V_2 \leftrightarrow V_b$, ... Using the adjacency lists of V_1 and V_a, we find that $V_9 \leftrightarrow V_f$. The adjacency lists of V_2 and V_b imply $V_{10} \leftrightarrow V_g$. The adjacency lists of V_3 and V_c imply: $V_7 \leftrightarrow V_l$, $V_8 \leftrightarrow V_i$, and $V_{10} \leftrightarrow V_g$,
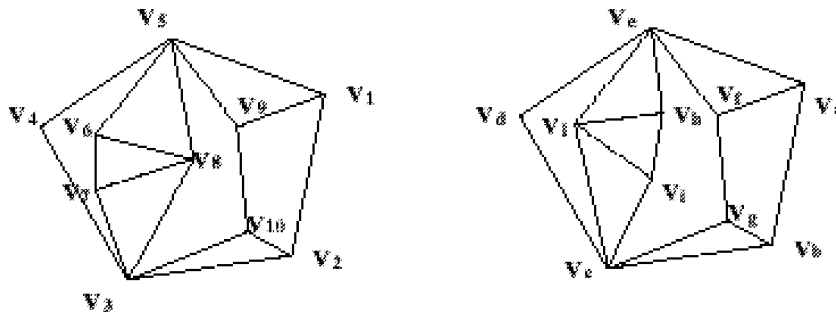
Fig. 10. Two line drawings whose TDs can be matched with a polynomial algorithm.

which had been already found. V_4 and V_d do not imply new correspondences. Finally, V_5 and V_e imply $V_9 \leftrightarrow V_f$, which was already found, but also $V_8 \leftrightarrow V_h$ and $V_6 \leftrightarrow V_l$, which are inconsistent with the previous pairings. Thus in this case the two TDs do not match; it can be easily verified that this also happens in the four remaining cases.

Using all the adjacency lists for each case takes $O(V)$ time, and since this must be done $m \leqslant V$ times, the algorithm for the case of the external polygon is $O(V^2)$.

Unfortunately, this algorithm does not apply to the case of internal aspects, where no external polygon is specified: in this case all possible isomorphism should be determined and checked. To be able to identify in the line drawings the polygons which are projections of the background would allow to apply the above algorithm also to many internal aspects. However, for multiply connected objects we could have worst case aspects where these polygons are $O(V)$, and can be paired in $O(V!)$ ways.

4. Topological recognition and similarity of polyhedra

We have seen that the topological comparison of aspects is mostly a fast process (for each aspect), but in the end it supplies approximate results. Topologically identifying a polyhedron S as one of the stored objects means that S is "similar" in some topological sense, not necessarily equal, to this object. For instance, it is clear that we are not able to topologically discriminate two different tetrahedra, since the topological AG is unique for all tetrahedra.

Of curse, as already observed, a topological match could be the first step of a more complex recognition process. In any case, our goal is to understand the theoretical capabilities and limits of the topological recognition. For this purpose, we address the following problem: *how similar are two polyhedra having topologically equivalent AGs?*

For answering this question, we state a reference condition for topological similarity of polyhedra, and compare this condition and the similarity induced by having the same AG.

The reference similarity condition is *isomorphism*:

Definition 1. Two polyhedra S and S' are said to be isomorphic if there exist one-to-one correspondences between edges and faces of S and edges and faces of S', and these correspondences are inclusion-preserving.
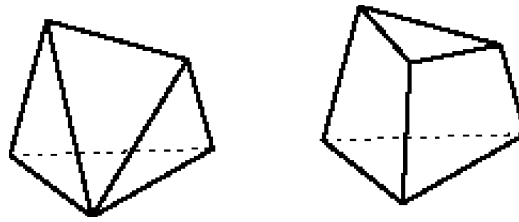
Fig. 11. Two pentahedra whose combinatorial types are different.

This definition, usually referred to topologically similar convex polytopes (see Gruenbaum [13, p. 38]), is also suitable for non-convex polyhedra. Informally, it means that for each face of *S* there is a corresponding face of *S'* with the same number of edges, and the corresponding faces are connected together in the same way. Only metric properties like lengths and angles can be different. Isomorphic polyhedra can be obtained one from the other by a continuous deformation which preserves the edges and the planarity of the faces. Isomorphism satisfies the natural idea of topologically similar polyhedra, especially in the convex case. In addition, it is the natural extension in 3-D of the 2-D topological equivalence between aspects as defined in the previous section. We can also produce an *a posteriori* reason for this choice. In the following, we will compare isomorphism and the similarity induced by the AG in a number of cases. Only in some of these cases this similarity is stronger than isomorphism: thus, it appears a similarity criterion "centered" on the problem we face.

Some further information on isomorphism of polyhedra could be interesting. Polyhedra with *n* faces can be grouped into *classes of isomorphism*, or *combinatorial types*. There is only one class for tetrahedra, but two classes for pentahedra (Fig. 11), and seven classes for esahedra.

The number of classes quickly increases with *n*: for instance, 257 different combinatorial types have been reported for $n = 8$ (see [13, p. 289]). Enumerating the isomorphism classes for any number of faces (or vertices) has been an open problem for about 150 years. Only recently an asymptotical answer has been found (see Bender [2]).

We shell see that the main theorem (Proposition 1) of the following section can be also expressed in terms of combinatorial types.

5. Perspective aspect graphs and isomorphism

In this section we compare, for *general* polyhedra, isomorphism and the similarity induced by having the same *perspective* aspect graph. We consider connected polyhedra, with vertices where at least three edges meet, and edges shared by two faces.

Some preliminary definitions are necessary:

Definition 2. Two aspects AS and AS' are equivalent if there exists a one-to-one correspondence between their vertices such that, by substituting in the topological description of one aspect the labels of the corresponding vertices of the other aspect, the two descriptions become equal.

Definition 3. Two aspect graphs G and G' are equivalent if there exist a one-to one correspondence between the vertices of G and G' such that:
  (i) the aspects of G are in a one-to-one correspondence with equivalent aspects of G';
 (ii) the transitions in G are in one-to-one correspondence with the transitions between the equivalent aspects in G'.

Definition 4. Two polyhedra are ag-equivalent if they have equivalent perspective aspect graphs.

We will show that ag-equivalence is a similarity condition *stronger* than isomorphism for general polyhedra. Although the proof that we have found is restricted to the class of general polyhedra with simply connected faces, this does not prevent these polyhedra to by multiply connected. Before deriving this result, some preliminaries are necessary.

Let us consider a face f of a polyhedron S and the closed sequence of edges which are the boundaries of f. It could be possible to find all the edges of this sequence in a unique aspect of the AG of S (for instance, this is always the case for convex faces). Otherwise, the various edges of the sequence can be traced in different aspects in the AG of S. Observe that there exists in the perspective AG of S at least one aspect containing completely any two consecutive edges of the sequence. A such aspect can be obtained from viewpoints sufficiently near to their common vertex. From this viewing region, three (at least) consecutive vertices of the sequence can be observed.

Consider now a closed ordered sequence of edges of the AG of a polyhedron S with simply connected faces, whose nodes are all *true* nodes and such that any two consecutive edges can be identified in at least one aspect. Let a such sequence of edges be indicated as an *e-sequence*. The following lemma gives a necessary and sufficient condition for an e-sequence to be the boundary of a face of S.

Proposition 1. *A necessary and sufficient condition for an e-sequence ES to be the boundary of a face f of S is that*: *in each aspect where two consecutive edges, i.e., three consecutive nodes V_{i-1}, V_i, V_{i+1} of the sequence appear, the edges incident at V_i, not belonging to ES lie all on the same side of the two edges with respect to the ordering assumed.*

Proof. The necessary part of the lemma is trivial. Let us consider a face f of S and assume for instance a clockwise ordering of its edges with respect to a viewpoint lying in the half-space containing the outward normal of the face. Let us consider an aspect where two consecutive edges of ES appear completely (Fig. 12). It is clear that the other edges incident at V_i must lie all on the same side of the e-sequence (on the left side in our case).

Observe that the area on the right side in Fig. 12, containing the point P, does not necessarily corresponds to the projection of f. Actually, the edges (V_{i-1}, V_i) and (V_i, V_{i+1}) could occlude another face f' (or the background).

Let us show now that the condition is sufficient. To this purpose, let us first consider two consecutive edges (V_{i-1}, V_i) and (V_i, V_{i+1}) of an e-sequence ES satisfying the condition, and show that they are edges of the same face. Let us suppose that this is false, i.e., that they do not belong to the same face. This contradicts the hypothesis that they belong to
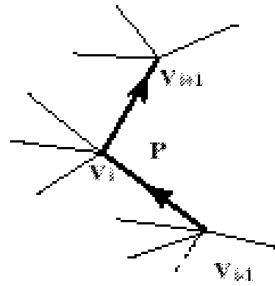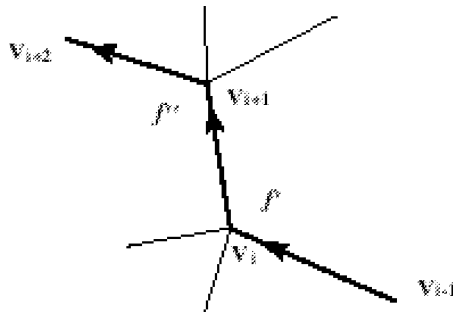
Fig. 12. Three nodes of an e-sequence.

Fig. 13. Three consecutive edges of an e-sequence do not satisfy the condition of Proposition 1 if they are not the boundaries of the same face.

an e-sequence that satisfies the condition of Proposition 1. In fact, it is always possible to find a viewing region near vertex V_i such that all the edges incident at V_i are visible, and therefore an aspect where there are edges on both sides of the e-sequence. Thus, we have shown by contradiction that any two consecutive edges of ES are edges of the same face.

Let us consider now three consecutive edges (V_{i-1}, V_i), (V_i, V_{i+1}), (V_{i+1}, V_{i+2}) of ES. We have that (V_{i-1}, V_i), (V_i, V_{i+1}) are edges of the same face, let it be f', and (V_i, V_{i+1}), (V_{i+1}, V_{i+2}) are edges of the same face, let it be f''. Could it be $f' \neq f''$? Obviously not, otherwise there would be incident edges on both sides of ES (see Fig. 13, possibly the composition of two adjacent aspects). From this, it follows that all edges of an e-sequence are the boundaries of the same face. $\square$

The necessary and sufficient condition of Proposition 1 will be indicated as Cond. 1. We can now state our main theorem:

Proposition 2. *If two general polyhedra S and S' with simply connected faces are ag-equivalent they also are isomorphic.*

Proof. We must prove a one-to-one, inclusion preserving, correspondence between edges and faces of S and edges and faces of S'. Let us consider a face f of S and the sequence of edges which form its boundary. This sequence of edges corresponds to a unique e-sequence
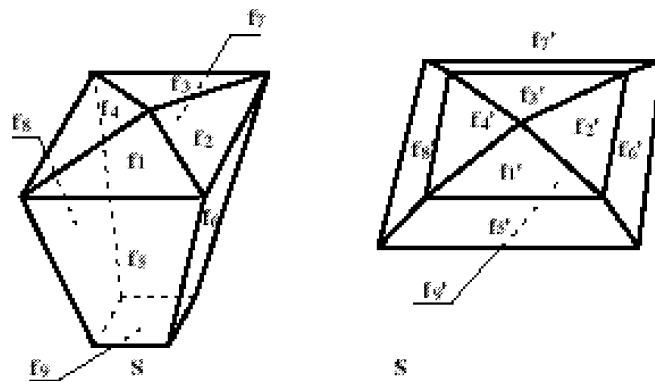
Fig. 14. Two isomorphic polyhedra which have different perspective aspect graphs.

ES in the AG of S. Since S and S' are ag-equivalent, the e-sequence ES is in one-to-one correspondence with an e-sequence ES$'$ of S'. We must now show that also ES$'$ identifies a face f' of S'. Proposition 1 guarantees that Cond. 1 holds for ES. The topological definition TD of equivalence of aspects preserves the order of the adjacency lists, so that Cond. 1 is true also for ES$'$. Proposition 1 again states that ES$'$ identifies a face f' of S'. Thus, ag-equivalence implies a one-to-one correspondence between the faces of S and S'. Since the faces are described in terms of edges (and the edges in terms of vertices), their one-to-one inclusion preserving correspondences follow. $\square$

Let us stress that isomorphism is necessary but *not sufficient* for ag-equivalence, or, in other terms, ag-equivalence is a stronger condition of topological similarity. For showing this, we present the following example. Let us consider the two isomorphic polyhedra S and S' in Fig. 14.

We can obtain S' by squeezing S and enlarging its bottom face f_9. In the AG of S' there is an aspect containing the 8 faces f_1', f_2', f_3', f_4', f_5', f_6', f_7', f_8'. A simple inspection of S shows that its AG does not contains any aspect with 8 faces.

We can reword the previous results as follows: *having the same perspective aspect graph divides polyhedra with n faces into a number of classes*: *each of these classes is a subset of one combinatorial class.*

As already observed, the similarity due to isomorphism is not so satisfactory for general polyhedra as for convex polyhedra. Consider for instance the isomorphic pairs (S_1, S_1') and (S_2, S_2') of Fig. 15(a) and (b). The point of these examples is that there are corresponding edges which are convex in one polyhedron and concave in the other isomorphic polyhedron (an edge shared by the faces f_i and f_j is said to be: (i) convex, if the angle between the outward normals of f_i and f_j is less than 180 degrees; (ii) concave, otherwise). In the following we will present another main property of ag-equivalence. It guarantees that pairs of isomorphic polyhedra like those in Fig. 15 are not ag-equivalent.

Proposition 3. *Two corresponding edges (V_i, V_j) and (V_i', V_j') of two ag-equivalent general polyhedra S and S' are both concave or both convex.*
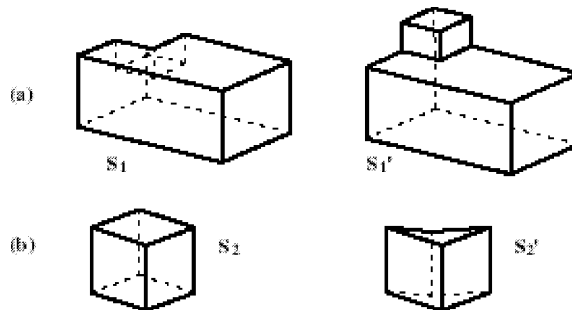
Fig. 15. Two pairs (S_1 and S'_1, S_2 and S'_2) of isomorphic polyhedra whose topological similarity is not very satisfactory. Such polyhedra cannot be ag-equivalent.
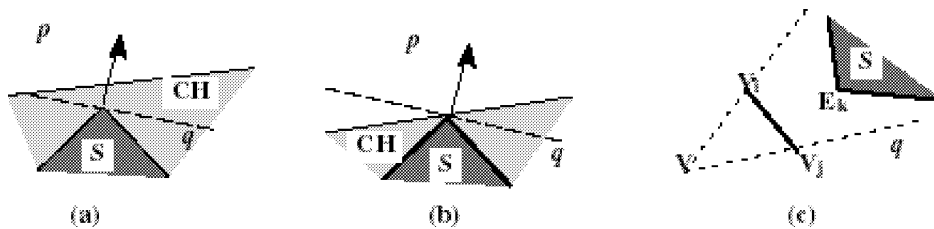

Fig. 16. The intersections of S and a plane p normal to a convex edge: (a) lying inside $\mathbf{CH}(S)$; (b) lying on the boundary of $\mathbf{CH}(S)$. In the plane q it is always possible to find viewpoints like V, such that the edge (V_i, V_j) occludes another edge (c), whose trace in q is marked E_k.

Proof. By contradiction, let us suppose that there is an edge (V_i, V_j) of S which is convex and whose corresponding edge (V'_i, V'_j) of S' is concave. Since the concave edge (V'_i, V'_j) cannot lie on the external contour of an aspect of S', also the convex edge (V_i, V_j) must not lie on the external contour of an aspect of S. This implies that (V_i, V_j) cannot lie on an edge of $\mathbf{CH}(S)$, the convex hull of S. Therefore, it must lie inside $\mathbf{CH}(S)$, or at most on a face of its boundary. For these two cases, Fig. 16 shows the intersection of S and a plane p normal to (V_i, V_j).

Consider a plane q which passes through (V_i, V_j) without entering S at (V_i, V_j) (see Fig. 16(a) and (b)). Plane q divides $\mathbb{R}^3$ into two half-spaces: consider the half-space marked by an arrow in Fig. 16. The intersection of this half-space and $\mathbf{CH}(S)$ must contain a part of S, otherwise, against the hypotheses: in case (a), (V_i, V_j) would lie on the boundary of $\mathbf{CH}(S)$; in case (b), (V_i, V_j) would be an edge of $\mathbf{CH}(S)$. In any case, since S is connected, plane q must intersect some faces of S (Fig. 16(c)). It is easy to understand that it is always possible to find a viewpoint V lying on q such that:

(i) (V_i, V_j) occludes one edge E_k of one of these faces, and
(ii) (V_i, V_j) is completely visible (Fig. 16(c)).

Thus, among the aspects of S there is at least one aspect where (V_i, V_j) partially occludes another edge. Obviously, this cannot occur in the aspects of S' containing the corresponding concave edge (V'_i, V'_j). This contradicts the hypothesis that S and S' are ag-equivalent. $\square$

6. Parallel aspect graphs and isomorphism

Many of the available algorithms compute the exact or approximate parallel AG. Parallel AGs are also used by most recognition algorithms. This choice is probably due to the complexity of the perspective AG and the algorithms for its computation.

Let us give the following definition:

Definition 5. Two polyhedra are pag-equivalent if they have equivalent parallel aspect graphs.

In this section we will compare isomorphism and the similarity induced by pag-equivalence for convex and general polyhedra.

6.1. Convex objects

It is easy to see that:

Proposition 4. *The pag-equivalence of two convex polyhedra S and S' is a sufficient condition for their isomorphism.*

It is clear that for any face of a convex polyhedron there is at least one parallel aspect containing a polygon whose edges represent all the edges of the face. Thus we do not need to introduce e-sequences for the proof, and pag-equivalence directly produces one-to-one correspondences between edges and faces of S and edges and faces of S'.

Actually, the equivalence of any pair of *complementary* parallel aspects is sufficient for isomorphism (an aspect is complementary to an other aspect of a convex polyhedron if it contains all the other faces of the polyhedron [23]).

Observe that isomorphism is not a sufficient condition for pag-equivalence. In fact, let us consider again the convex objects in Fig. 14. They are isomorphic but not pag-equivalent, since the aspect with 8 faces, relative to S' but not to S, is a parallel aspect. Thus, for convex polyhedra, pag-equivalence is a similarity condition stronger than isomorphism.

It can also be shown that pag-equivalence is weaker than ag-equivalence, as it can be guessed considering that pag-equivalence is the equivalence of only a subgraph of the perspective AG. An example will show that there are convex objects with the same parallel aspect graph, but with different perspective aspect graph. Consider the isomorphic objects S and S' in Fig. 17.

For each object, three views along three ortogonal directions and a trimetric view are shown in the figure. The height of the faces f_1, f_2, and f_3 of S is greater than the height of the corresponding faces of S'. It can be easily verified (we omit the details) that S and S' have equivalent parallel AGs. Let us consider the tetrahedral viewing region **VR** of the perspective AG of S, highlighted in Fig. 18. It is bounded by the planes supporting the faces f_5, f_8, f_2, f_3 of S.

Only the faces f_1, f_2 and f_3 are visible from **VR**, and thus the aspect A* relative to this region contains three polygons each with four edges. An inspection of S' shows that no similar viewing region, able to observe f'_1, f'_2 and f'_3 exists. Since no other face has four
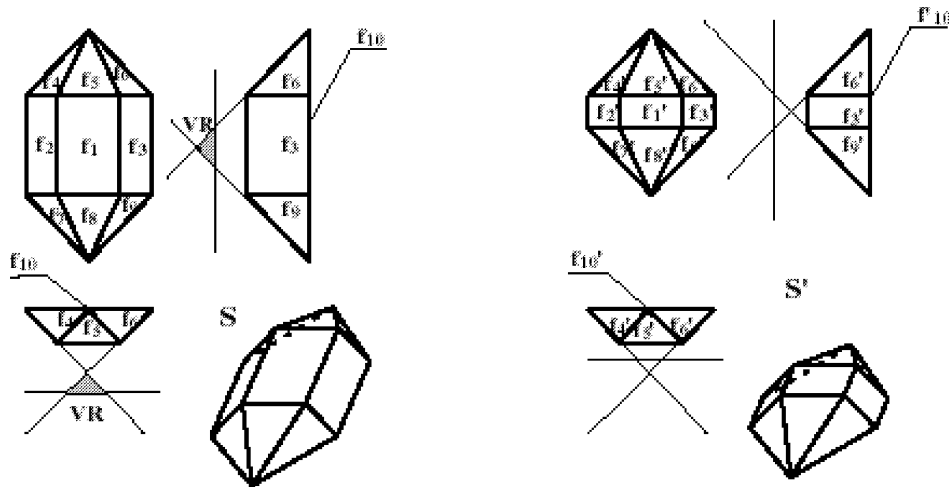
Fig. 17. The convex polyhedra S and S' have equivalent parallel aspect graphs, but different perspective aspect graphs. The highlighted viewing region **VR** supplies an aspect with three faces each with four edges, which does not exists in the perspective AG of S'.
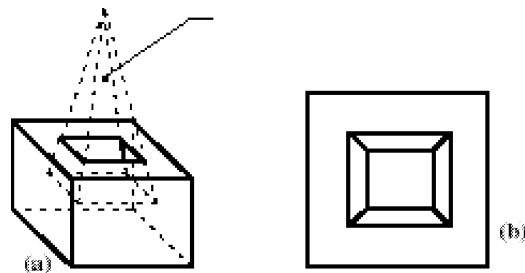

Fig. 18. A concave polyhedron (a), and the perspective aspect observed from **VR** (b), which is not contained by any parallel aspect.

edges, it is also clear that no aspect equivalent to A* can be found in the AG of S'. Thus S and S' are pag-equivalent but not ag-equivalent.

Summarizing, pag-equivalence of convex polyhedra, although weaker than ag-equivalence, is stronger than isomorphism, and therefore the parallel AG can be used for discriminating non-isomorphic convex objects.

It must be stressed that using the parallel AG for object recognition puts limits on the position of the vantage points. Actually, they should belong to the unbounded viewing zones only. However, the bounded viewing zones, which supply perspective aspects, could extend at any distance from the object. For instance, we can move region **VR** in Fig. 17 arbitrarily away from S by conveniently reducing the two angles between the pairs of faces f_5, f_8 and f_2, f_3.

In order to cope with this problem, let us observe that, for convex polyhedra, any perspective aspect is contained in at least one parallel aspect [23]. We say that an aspect A

contains another aspect A′ if A′ can be obtained by deleting successively from A a number of *external* polygons, i.e., with at least an edge on the outer boundary of the aspect. On the basis of this property, we could deal with viewpoints inside bounded zones by attempting to match perspective aspects with sub-aspects of the parallel aspects. However, this is a partial remedy, since to be a sub-aspect of the parallel AG is necessary, but in general not sufficient, for the aspect to belong to the perspective AG of the object, as shown by the example of Fig. 17.

6.2. General polyhedra

We have the obvious statement:

Proposition 5. *For non-convex polyhedra, pag-equivalence is not sufficient for isomorphism.*

It is trivial to find examples of non-isomorphic concave polyhedra with the same parallel AG. In these cases, it is clear that the polyhedra are different (i.e., non-isomorphic) at "internal" parts only. The recent introduction of the concept of *Internal Visual Hull* (**IVH**(S)) of an object S allows a more precise formulation of this statement (see [24]). Only the edges which belong to the surface of **IVH**(S) can appear in parallel aspects; thus edges which do not satisfy this condition are irrelevant for pag-equivalence and for recognition from parallel aspects.

Obviously, as for convex polyhedra, using the parallel AG for object recognition puts limits on the position of the viewpoints. The restriction is more severe in this case, since perspective aspects of non-convex polyhedra could be not contained by parallel aspects. In fact, let us consider the concave polyhedron and the pyramidal viewing region **VR**, which can extend at any distance from the object, shown in Fig. 18(a).

In Fig. 18(b) it is shown a line drawing representing the perspective aspect relative to **VR**, which is not contained in any parallel aspect.

In conclusion, for discriminating concave polyhedra the parallel AG appears not as effective as it is for convex objects.

7. Reduced aspect graph, aspect set and isomorphism

Identifying aspects and visual events from 2-D images requires monitoring an object at sufficiently short time intervals. This is necessary for understanding when a visual event takes place, and matching the corresponding vertices in the images before and after the transition.

In practice, a set of "snapshots" of the unknown object supplies a sequence of aspects, but we could unable to understand if one aspect is adjacent to the previous one, or to identify the corresponding vertices. In these situations, the full information provided by the AG is redundant, and we can perform the topological identification using reduced stored information, since several topologically equivalent aspects may appear in an AG.

7.1. The reduced aspect graph and the aspect set

First, let us assume that we are able to understand whether two aspects are adjacent, but unable to, or not interested in identifying the corresponding vertices of two adjacent aspects. This suggests the idea of *reduced aspect graph* (*RAG*) of a polyhedron [25]. The RAG can be obtained from the full AG by merging together all the nodes labeled with topologically equivalent aspects. The arcs incident at nodes merged become arcs incident at the unique resulting node. If two or more arcs joining a pair of nodes are obtained, all of them but one should be deleted. Each aspect of the RAG represents a category of topologically equivalent aspects of AG. In each aspect, the labels of the vertices are not in correspondence with those of the other aspects. Merging aspects together had been proposed also in [31,33].

The RAG has the following main property: *any sequence of adjacent aspects which can be found in the full AG of a polyhedron can also be found in its RAG.*

This property, together with the reduced size, makes the RAG a candidate to effectively support topological recognition from multiple images. As already observed in Section 3, the topological comparison of line drawings and stored aspects is also simplified. In fact in this case it is not necessary to find all the correspondences which satisfy TD, as required for the regular AG. However, two polyhedra having equivalent RAGs could have different AGs (a simple convex example can be found in [25]), and therefore in general the AG has greater capabilities than the RAG for discriminating objects.

If we are not interested in, or not able to understand whether two aspects of the unknown object are adjacent, a further reduction of the stored information can take place. By deleting from the RAG all the arcs, we obtain the *aspect set* (*AS*), that is the set of all the topologically different aspects of the polyhedron.

In Fig. 19, the RAGs of a tetrahedron (a), a pentahedron (b), a parallelepiped (c), are presented by means of a representative line drawing for each aspect.

The number n_{as} of aspects in the RAG and in AS could be much smaller than the number of aspects in the AG. For instance, with reference to the objects in Fig. 19, the pentahedron has 25 aspects in the perspective AG and 8 in the aspect set; these numbers are 26 and 3 for the cube and 14 and 3 for the tetrahedron.
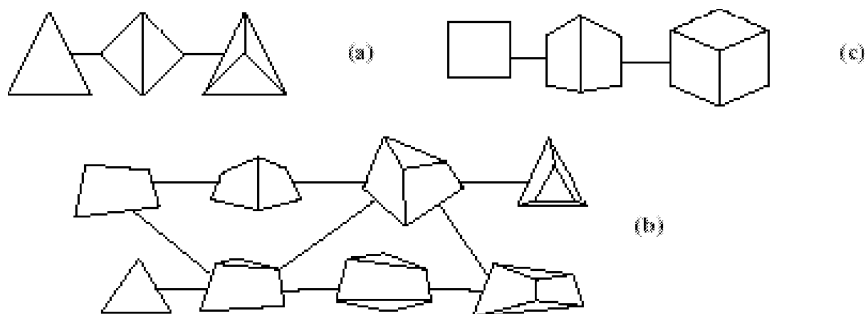

Fig. 19. The reduced aspect graphs of: a tetrahedron (a), a pentahedron (b), a parallelepipedon (c).

For the numbers of aspects of general convex polyhedra and polyhedra with parallel faces, the reader is referred to [13, p. 393] and [23].

Finding average values and tight upper bounds for n_{as} does not appear a simple task, even for convex polyhedra. However, it is easy to find a lower bound (we omit the details): $\lfloor n/2 \rfloor \leqslant n_{as}$ for convex polyhedra with n faces. This is a tight bound, for instance, for a parallelepiped.

7.2. Reduced aspect graph, aspect set and isomorphism

Which are the relations between *rag-equivalence*, *as-equivalence* (we omit for brevity these obvious equivalence definitions) and isomorphism? We have found the following result:

Proposition 6. *Both rag-equivalence and as-equivalence are not sufficient conditions for isomorphism of general polyhedra.*

Proof. Let us consider the polyhedron C in Fig. 20. It is a cube traversed by a hole with a square section and 4 sharp "bends". The shape of the hole is shown in Fig. 20(b).

Observe that the geometry of the hole is such that viewpoints inside the hole cannot see more than three "bends" at a time. Let us consider another non-isomorphic cube C', containing a similar hole but with a greater number of "bends". Provided that also in this hole it be impossible to see more than three "bends" at a time, it is clear that C' has exactly the same aspect set and the same reduced aspect graph than C.   □

Let us show now that as-equivalence (and rag-equivalence) are not necessary conditions for isomorphism. Consider again the isomorphic polyhedra in Fig. 14. We have seen that they have different AGs. Also their RAGs and ASs are different, since an aspect with 8 faces is contained in the AG, and therefore in the AS, of S but not in the AG, an therefore not in the AS, of S'.

Concluding, for general polyhedra, isomorphism cannot be related to rag-equivalence and as-equivalence.
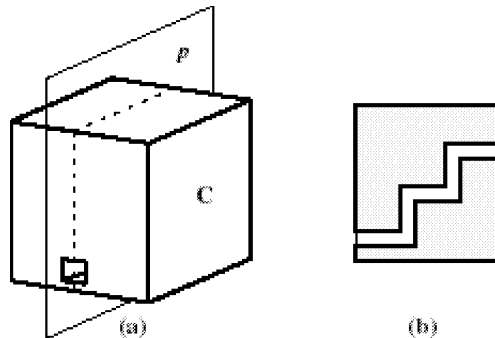

Fig. 20. A cube C traversed by a square section hole (a), and its section made by plane p.

8. Conclusions

This paper provides a theoretical frame for using the AG for topological recognition from multiple views.

First, we have developed a definition of aspect of polyhedra suitable for topological recognition, and discussed the computational complexity of comparing aspects.

For understanding limits and capabilities of topological recognition, we have also addressed the problem of understanding how similar are two polyhedra with the same (topological) AG. The similarity implied by the perspective AG and the parallel AG have been compared with isomorphism for general and convex polyhedra.

The smaller size, as well as the practical problems in observing visual events, make the reduced (perspective) AG and the (perspective) AS interesting tools for topological recognition. The topological similarity that they imply has also been compared with isomorphism. The results obtained are summarized in Table 1.

It is worth noting that general polyhedra with the same perspective AG have been found not only to be isomorphic, but also to have two corresponding edges both concave or convex. Therefore ag-equivalence is a rather strong similarity condition.

On the other hand, parallel AGs are guaranteed to be able to discriminate non-isomorphic objects only in the convex case. Even objects with the same reduced aspect graph and aspect set could be non-isomorphic. However, in practical situations, with a limited set of not tricky objects, these theoretical limits do not prevent a fruitful use of the parallel AG, the RAG and the AS.

A number of theoretical questions remains open. Although significant necessary conditions for ag-equivalence have been given (isomorphism and convexity or concavity for corresponding edges), a closer characterization of ag-equivalence could be obtained if necessary and sufficient conditions were found. Another subject deserving further investigation is the power of discriminating polyhedra of both RAG and AS. For showing that the rag-equivalence and as-equivalence do not imply isomorphism for general polyhedra, we have used a rather extreme example. These equivalences can be conjectured to imply isomorphism for less unusual objects. For convex polyhedra we have been unable to prove that as-equivalence implies isomorphism, or to find a counter-example. However, we succeeded in proving that two as-equivalent convex polyhedra have the same number of faces, and we conjecture that they are isomorphic.

Table 1

Similarity due to:	Perspective AG	Parallel AG		Reduced AG (perspective)	Aspect set (perspective)
Type of polyhedron	General*	General	Convex	General	General
Relation with isomorphism	$\geqslant$	$\leqslant$	$\geqslant$	$\leqslant$	$\leqslant$

$\geqslant$ implies isomorphism;
$\leqslant$ does not imply isomorphism;
* proof given for simply connected faces.

References

[1] M. Behzad, G. Chartrand, Introduction to the Theory of Graphs, Allyn and Bacon, Boston, MA, 1971.

[2] E. Bender, The number of three-dimensional convex polyhedra, Amer. Math. Monthly 94 (1987) 7–21.

[3] K. Bowyer, C. Dyer, Aspect graphs: An introduction and survey of recent results, Internat. J. Imaging Systems and Technology 2 (1990) 315–328.

[4] K. Bowyer et al., Computing the generalized aspect graph for objects with moving parts, IEEE Trans. Patt. Anal. Machine Intell. 15 (6) (1993) 605–610.

[5] N. Chiba et al., A linear algorithm for embedding planar graphs using PQ-trees, J. Comput. System Sci. 30 (1985) 54–76.

[6] S. Dickinson, A. Pentland, A. Rosenfeld, 3-D shape recovery using distributed aspect matching, IEEE Trans. Patt. Anal. Machine Intell. 14 (2) (1992) 174–198.

[7] D. Eggert, K. Bowyer, C. Dyer, H. Christensen, D. Goldgof, The scale space aspect graph, in: Proc. IEEE Conference on Computer Vision and Pattern Recognition, 1992, pp. 335–340.

[8] D. Eggert, K. Bowyer, Computing the perspective projection aspect graph of solids of revolution, IEEE Trans. Patt. Anal. Machine Intell. 15 (2) (1993) 109–128.

[9] O. Faugeras, J. Mundy, N. Ahuja, C. Dyer, A. Pentland, R. Jain, K. Ikeuchi, K. Bowyer, Panel theme: Why aspects graphs are not (yet) practical for computer vision, CVGIP: Image Understanding 55 (2) (1992) 212–218.

[10] G. Fekete, L.S. Davis, Property spheres: A new representation for 3D recognition, in: Proc. IEEE Workshop on Computer Vision, Representation and Control, 1984, pp. 192–201.

[11] H. Gazit, J.H. Reif, A randomized parallel algorithm for planar graph isomorphism, in: Proc. 2nd Ann. ACM Symposium on Parallel Algorithms and Architectures, New York, 1990, pp. 210–219.

[12] Z. Gigus, J.F. Canny, R. Seidel, Efficiently computing and representing aspect graphs of polyhedral objects, IEEE Trans. Patt. Anal. Machine Intell. 13 (6) (1991) 442–451.

[13] B. Gruenbaum, Convex Polytopes, Wiley, New York, 1967.

[14] J. Hopcroft, J.K. Wong, Linear time algorithms for isomorphism of planar graph, in: Proc. Sixth Ann. ACM Symposium on Theory of Computing, 1974, pp. 172–184.

[15] K. Ikeuchi, Generating an interpretation tree from a CAD model for 3d-object recognition in bin-picking tasks, Internat. J. Comput. Vision (1987) 145–165.

[16] J. Jaja, S. Rao Kosaraju, Parallel algorithms for planar graph isomorphism and related problems, IEEE Trans. Circuits and Systems 35 (1988) 304–311.

[17] R. Jayakumar et al., Planar embedding: Linear time algorithms for vertex placement and edge ordering, IEEE Trans. Circuits and Systems 35 (1988) 334–344.

[18] L.M. Kirousis, Effectively labeling planar projections of polyhedra, IEEE Trans. Patt. Anal. Machine Intell. 12 (2) (1990) 123–130.

[19] J.J. Koenderink, A.J. van Doorn, The internal representation of solid shapes with respect to vision, Biol. Cybern. 32 (1979) 211–216.

[20] M.R. Korn, C.R. Dyer, 3-D multiview object representations for model-based object recognition, Pattern Recognition 20 (1987) 91–103.

[21] D. Kriegman, J. Ponce, Computing exact aspect graph of curved object: Solid of revolution, Internat. J. Comput. Vision 5 (2) (1990) 119–135.

[22] D. Kriegman, J. Ponce, Computing exact aspect graphs of curved objects: Parametric surfaces, Technical Report UIUC-DCS-R-90–1579, Comput. Sci. Dept., Univ. of Illinois at Urbana-Champaign, IL, 1990.

[23] A. Laurentini, Equivalent aspects of convex polyhedra, in: Proc. Bilkent. I. Conf. New Trends Comm., Control Signal Proc., Elsevier, Amsterdam, 1990, pp. 1750–1756.

[24] A. Laurentini, The visual hull concept for silhouette-based image understanding, IEEE Trans. Patt. Anal. Machine Intell. 16 (2) (1994) 150–162.

[25] A. Laurentini, Introducing the reduced aspect graph, Pattern Recognition Letters 16 (1995) 43–48.

[26] U. Manber, Introduction to Algorithms, Addison-Wesley, Reading, MA, 1989.

[27] S. Petitjean, J. Ponce, D.J. Kriegman, Computing exact aspect graph's of curved objects: Algebraic surfaces, Int. J. Comput. Vision 9 (3) (1992) 231–255.

[28] W.H. Plantinga, C.R. Dyer, Visibility, occlusion and the aspect graph, Internat. J. Comput. Vision 5 (2) (1990) 137–160.

[29] M. Seibert, A.M. Waxman, Adaptive 3-D object recognition from multiple views, IEEE Trans. Patt. Anal. Machine Intell. 14 (2) (1992) 107–124.

[30] I. Shimshoni, J. Ponce, Finite-resolution aspect graphs of polyhedral objects, IEEE Trans. Patt. Anal. Machine Intell. 19 (4) (1997) 315–326.

[31] J. Stewman, L. Stark, K. Bowyer, Restructuring aspect graphs into aspect- and cell-equivalence classes for use in computer vision, in: Graph Theoretical Concept in Computer Science, Springer, New York, 1988, pp. 230–241.

[32] S.A. Underwood, C.L. Coates Jr., Visual learning from multiple views, IEEE Trans. Comput. 24 (1975) 651–661.

[33] R. Wang, H. Freeman, Object recognition based on characteristic view classes, in: Proc. 10th Internat. Conference on Pattern Recognition, 1990, pp. 8–13.