

Part segmentation of objects in real images

Guillaume-Alexandre Bilodeau, Robert Bergevin^{*}

Computer Vision and Systems Laboratory, Pavillon Adrien-Pouliot, Université Laval,

Ste-Foy (Québec), G1K 7P4, Canada

Accepted for publication in Pattern Recognition

Keywords: Part segmentation; Perceptual grouping; Image query; Volumetric primitives; Geons; Constant-curvature contour primitives;

^{*} Corresponding author. Tel: +1 418 656 2131 x5173 fax: +1 418 656 3594
E-mail address: bergevin@gel.ulaval.ca (R.Bergevin)

Abstract

This paper addresses one major aspect of a project aimed at the design of an image database query engine, where the images are searched at the 3D object-level. This approach is a novelty since image databases are usually searched by comparing colours, textures and 2D shapes of regions in the images. The main topic and contribution of this paper is a new method for object part segmentation, from constant-curvature contour primitives, that is suited to process images of objects in real scenes. The proposed method groups constant-curvature primitives using intermediate-level geometric relationships and object outline information. A detailed description of the method is presented along with validating experiments.

Keywords: Part segmentation; Perceptual grouping; Image query; Volumetric primitives; Geons; Constant-curvature contour primitives;

1. Introduction

The comparison of 3D objects in 2D images using efficient and reliable algorithms is a problem that remains unresolved in computer vision. A similar problem is the identification of an object in an image. A value of similarity is obtained as the result of the comparison of two objects, whereas one or more identifiers are produced when an object is identified. The images are processed with common algorithms, but the results obtained are interpreted differently. The resolution of these two problems is of high interest as it permits the development of autonomous robots and efficient image database query engines.

Our work aims precisely at developing robust algorithms to model and compare 3D objects in 2D images in the context of an image database query engine. In this paper, one specific aspect is presented. It is the segmentation of a 3D object into 2D parts. In the context of this work, 2D

parts are defined as regions delimited by groups of circular arcs and straight-line segments (constant-curvature primitives), which can later on be interpreted as the projections in the plane of simple geon-like volumetric primitives [1]. Thus, after interpretation, the segmentation we are looking for will produce a spatial arrangement of volumetric primitives that resemble the spatial arrangement of the volumetric primitives making out the depicted 3D object. This will ensure at least partial view independence for the object models.

The main contribution of this paper is a method for the detection of parts by grouping two constant-curvature primitives (CCPs) using the outline to reduce the number of possible groups. We will show that our approach is suited to process 3D objects in 2D images of real scenes by several validating experiments.

The paper is structured as followed. Section 2 gives an overview of the application. Section 3 provides a review of the literature and our basic strategy to part segmentation. Section 4 and 5 present our method in detail. Section 6 gives an analysis of results obtained and Section 7 concludes the paper.

2. Overview of the envisioned application

The application under development aims at querying an image database of manufactured objects, which are interpretable as arrangements of simple volumetric primitives. Images in the database are from real scenes with one main object in the foreground that must be detectable in the image (see [2]).

Figure 1

Figure 1 shows an overview of the image database query engine under development. The shaded region represents the four modules required to add an image or query the database. The database is composed of various 2D images of 3D objects, and their associated models. To query or add an image to the database, the user gives as input an example 2D image or a sketch of the 3D object. The image is first processed to obtain contours of linked local intensity edges that are segmented to produce a map of constant-curvature primitives (CCPs) [3]. An initial grouping of the CCPs produces the outline of the object (Object detection) [2]. The CCP map is then processed further to obtain parts using the extracted outline (Part segmentation). These parts are labelled based on the possible volumetric primitives that may project onto them (Object modeling). Parts are interpreted as volumetric primitives since the aspect of a projected 3D object may change significantly for different viewpoints (see Figure 11). The object modeling module also computes the spatial relationships between parts. Finally, the constructed model is compared with the models in the database (Model matching). If similar models are in the database, the corresponding 2D images are shown to the user. If not, the newly built model and its corresponding image may be added in the database.

Therefore, the general goal of the image database query engine is showing to the user the images in the database that resemble the most the query image or sketch, and at the same time limiting the number of false positives. The obtained images will be classified from the most to the least similar image based on the score obtained during matching. The topic of this paper is the second module: Part segmentation. The following section reviews the literature and describes our approach to this problem.

3. Possible approaches to part segmentation

Five types of approaches have been proposed to solve the part segmentation problem. They are the approaches based on: i) symmetry axes of the parts of the object, ii) convex dominant points on the boundary of the object, iii) minimum description length criterion, iv) closed paths of primitives, and v) perceptual grouping and geometric relationships. We review them briefly and explain why none is appropriate as such for our problem.

3.1. Approaches based on symmetry axes

Several works ([4],[5],[6],[7]) have used symmetry axes to segment objects into parts. After edge detection, the object symmetry axes are detected. The axes are usually computed by taking the midpoint between two edge pixels, and by linking these midpoints. The result is a collection of axes in which some are relevant for part segmentation and some are not. Rules are then used to select the axes that possibly describe parts. The boundary that has given rise to a selected axis corresponds to the boundary of a part.

These approaches work well for objects that are lightly textured. However, in our application, heavily textured object may be processed. For heavily textured objects, these approaches give rise to a very large collection of axes (the textures give rise to axes) from which it is very difficult to select the relevant ones. Hence, for our application, these approaches are not appropriate.

3.2. Approaches based on convex dominant points

These approaches are illustrated by the work of Bennamoun and Boashash [8]. The basic strategy in these approaches is to segment the object based on the convex dominant points on its boundary. Convex dominant points are locations on the boundary where there is a significant

orientation change between consecutive groups of boundary elements. These approaches suppose that a pair of convex dominant points defines a joint between two parts. Hence, after edge detection, the boundary is first scanned to locate these points. Then, they are paired to find the joints between the parts. Once the joints are found, the parts are defined by the boundary lying between joints.

These approaches are not suited for our application, because some convex dominant points may be absent due to gaps between edges. In addition, objects, which can self-occlude in some viewpoints, cannot be segmented invariably into the same group of parts. Some convex dominant points can appear and disappear according to the viewpoint and change the pairing and the segmentation.

3.3. Approaches based on a minimum description length criterion

These approaches group edge points by template matching using a minimum description length (MDL) criterion. In the work of Pihu and Fisher [9], ellipsoid templates acting as parts are fitted on the edge map. The length of description may be calculated in the following way for an edge map. The associated length is zero for parts of the ellipsoid that are on boundary pixels. For parts of the ellipsoid that do not fit any boundary pixel, the length is increased by one for every pixel that does not fit. Therefore, the smaller the description length, the better the fit. In these approaches, the MDL criterion is optimized globally for the entire image to extract the best combination of parts.

In many cases, the parts obtained with these approaches do not capture the actual shape and structure of the objects. Hence, a single ellipsoid may be fitted to more than one projection of volumetric primitives and it may not fit well on volumetric primitives with concavities. This is

caused by the lack of intermediate structures between edge points and parts. Therefore, these approaches do not have the generality and invariance needed for our application.

3.4. Approaches based on closed paths of primitives

These approaches attempt to build closed paths of primitives (e.g.: circular arcs or straight-line segments). Therefore, the 2D image is first transformed into a graph making explicit geometric relationships among extracted primitives. Then, paths are constructed in such a way that they should correspond to parts. Rules are defined for validating the addition of a new primitive in a given path. The work of Jacot-Descombes and Pun [10] implements this basic approach. Their graph of the primitives (the nodes) and of the junctions (the arcs) of the object is constructed based on proximity. From this graph, all the possible closed paths are found, and those that do not contain other closed paths are considered as parts. The work of Jacobs [11] differs by the use of a convexity criterion and a saliency criterion on the closed paths.

These approaches are not suited for our application, since their grouping criteria do not consider the 3D context, for instance occlusions and self-occlusions of parts, which is essential in processing images of real scenes.

3.5. Approaches based on perceptual grouping and geometric relationships

These last approaches group constant-curvature primitives (CCPs) extracted from a 2D image based on geometric relationships. They are inspired by works on perceptual grouping ([1],[12]). The spatial arrangements of the CCPs are studied, and those that fit some criteria based on geometric relationships are considered as parts. Examples are PARVO [13] and the system developed by Dickinson and al. [14]. PARVO uses symmetry and junctions to group CCPs into faces and then parts. Junctions give the clues about the structure of the object and ensure robust

grouping. Dickinson system's groups CCPs in many stages. First, arrangements of CCPs are searched. Then, these arrangements of CCPs are grouped into faces, the faces are grouped into aspects, and the aspects are grouped into geons. All groups that are built must belong to a predefined catalogue of multi-level arrangements (CCPs, faces, aspects, geons). Probabilities are associated to each grouping hypothesis.

Images of real scenes are processed in our application, and some faces may not be extracted. In PARVO, a single missing face always inhibits the proper extraction of the part. Our new method removes this limitation.

The problem with the approach of Dickinson is that in practice, it is not possible to predict all specific arrangements of CCPs, faces and aspects that may be found in real images. This is because part boundaries are often fragmented due to noise and occlusions.

Although the existing implementations of these approaches are not suited for our application, the principle behind them can be adapted to fit its needs. The following section explains why, and how it can be done.

4. Principle of our method

As mentioned previously, our application needs to automatically process images depicting objects in real scenes. Grouping constant-curvature primitives (CCPs) by geometric relationships is valid in this context if the CCPs are grouped in such a way that the influence of texture and noise is reduced. This can be done using structural information about the object. A convenient feature for this purpose is the outline of the object. In our system, it is available as a set of CCPs. Although it conveys only limited information on the object structure, it does not restrict the types

of arrangements of CCPs that can be processed. Furthermore, its simplicity allows a user to make a query by sketching only the outline of an object or each of its constitutive parts. It is to be noted that the outline was also used in PARVO but in a different fashion than for our new method.

The principle of our method for part segmentation is to group CCPs into pairs using a criterion based on intermediate-level geometric relationships, while accounting for structural constraints from the outline. These constraints are associated to the way parts arise from protrusions and indentations in a multipart object. They are related, but simpler than the minimum-cut rule of Hoffman [15]. Subsequently, the complete boundary of the part is detected by making paths with CCPs of the object to link the two grouped CCPs. The shape of the parts extracted is only limited by the structure of the object (obtained via the outline) and by the geometric relationships between the two CCPs grouped. Note that the outline is extracted using the method of Bilodeau and Bergevin [16] and the CCPs are extracted using MAGNO ([17]).

This method has the advantage of being able to extract any type of parts that are made of two compatible CCPs (as defined by the geometric criterion), along with their detailed shape.

Therefore, the parts extracted may be non-convex. Furthermore, since the parts are made of two compatible CCPs, they can easily be interpreted as simple volumetric primitives. Using the definition of generalized cylinders, the two compatible CCPs are interpreted as the non-sectional sides of the projected envelope of the volume swept by the section, and the CCPs used to link the two compatible CCPs describe the projected section. In addition, the use of the outline reduces the complexity of the search for two compatible CCPs.

5. Details of part segmentation process

Let us now detail the part segmentation process. The solution space of our part segmentation method is the set of all possible CCPs pairs. The goal of our method is to select a subset of pairs of CCPs among these, which correspond to parts. This is done using a criterion that increases the probability that the selected pairs of CCPs correspond to the projected envelopes of the simple volumetric primitives composing the object. Once two compatible CCPs of a part are found and validated, the boundary of the part is completed by adding the CCPs that correspond to the section.

Since not all parts are on the outline of the object, part segmentation is accomplished in two stages. Stage 1 detects parts with the help of the outline and intermediate-level geometric relationships. Parts detected at this stage are more reliable than at the following stage because of the structural constraints from the outline. Stage 2 relies solely on intermediate-level geometric relationships. Only CCPs that are on the outline (outline CCPs) or inside the area bounded by the outline (interior CCPs) are considered for grouping. The geometric criterion used for grouping is described in the next section, and then the two stages are detailed.

5.1. Geometric criterion

A unique criterion based on geometric relationships (CBGR) is used for both stages. This grouping criterion is based on five intermediate-level geometric relationships between two constant-curvature primitives (CCPs). They are the parallelism (P), the distance between the CCPs (D), the similarity in length (SL), the similarity in type (ST) and the overlap (O). These relationships are considered as intermediate-level, because they are computed at the CCP level. Each relationship is normalized and the grouping criterion is formulated as follows:

$$CBGR = \frac{P + D + SL + ST + O}{5} \quad (\text{EQ. 1})$$

The next subsections show how the five intermediate-level geometric relationships are defined.

5.1.1. Parallelism (P)

Parallelism is in general a non-accidental property of the projected envelopes of simple volumetric primitives. This geometric relationship is the component of the grouping criterion used to support the grouping of two CCPs that are as close as possible as being parallel. The parallelism is computed in a similar fashion for both straight-line segments and circular arcs. Hence, the parallelism is calculated by approximating the circular arcs by straight-line segments joining their endpoints. Defining the straight-line segments as vectors pointing in an arbitrary chosen direction, parallelism is given by,

$$P = \left| \vec{a} \cdot \vec{b} \right| \quad (\text{EQ. 2})$$

where \vec{a} and \vec{b} are two vectors representing straight-line segments. The values obtained for P are between 0 and 1.

5.1.2. Distance between the CCPs (D)

The distance between the CCPs is used to ensure that the two compatible CCPs of a part are as close as possible. This geometric relationship is part of the grouping criterion because two close CCPs are more likely to define the boundary of the projected envelope of a single simple volumetric primitive. The distance between two CCPs is defined by,

$$D = 1 - \frac{\|mpCCP_1 - mpCCP_2\|}{\max(\dim x, \dim y)} \quad (\text{EQ. 3})$$

where $mpCCP_1$ and $mpCCP_2$ are the midpoints of the two CCPs, and $\dim x$ and $\dim y$ are the number of columns and rows in the image. This value is between 0 and 1.

5.1.3. Similarity in length (SL)

The similarity in length is used to support the grouping of two compatible CCPs that have about the same length. This is because in general, the two CCPs defining the projected envelope of simple volumetric primitives have similar lengths. For the similarity in length, the circular arcs are approximated by straight-line segments in a similar fashion as for the parallelism. The actual length of the circular arcs is not used since the similarity in the distance between the endpoints is the property of interest. This is because we are interested in obtaining enclosed regions that are as symmetrical as possible. The similarity in length is calculated as follows:

$$SL = \frac{\min(l_1, l_2)}{\max(l_1, l_2)} \quad (\text{EQ. 4})$$

where l_1 and l_2 are the length of the CCPs. This value is between 0 and 1.

5.1.4. Similarity in type (ST)

The similarity in type favours the grouping of CCPs of the same type (circular arcs or straight-line segments). This is justified by the fact that, in general, simple volumetric primitives have their envelope defined by two CCPs of the same type. The similarity in type is computed by comparing the type of the two CCPs. The similarity in type is given by,

$$ST = \begin{cases} 1 & \text{if } Type(CCP_1) = Type(CCP_2) \\ 0 & \text{if } Type(CCP_1) \neq Type(CCP_2) \end{cases} \quad (\text{EQ. 5})$$

where $Type()$ is an operator that returns the type of a CCP, and CCP_1 and CCP_2 are the two tested CCPs. This value is between 0 and 1.

5.1.5. Overlap (O)

For simple volumetric primitives, the two CCPs defining their envelopes are in general overlapping. The overlap relationship component of the grouping criterion supports the pairs of CCPs that are overlapping. The overlap (O) for a pair of straight-line segments is defined as the length of the projection of CCP_2 on CCP_1 (Lo_1) and the length of the projection of CCP_1 on CCP_2 (Lo_2) divided by the sum of the length (L_1) of CCP_1 and of the length (L_2) of CCP_2 . That is:

$$O = \frac{Lo_1 + Lo_2}{L_1 + L_2} \quad (\text{EQ. 6})$$

This value is also between 0 and 1. For circular arcs, the overlap is calculated by approximating them by straight-line segments as for the parallelism.

As an example of results for the grouping criterion, the pair of dark CCPs in Figure 2 (Group X) has a grouping criterion value of 0.98, and the values of its components are $P=0.99$, $D=0.94$, $SL=0.99$, $ST=1$, and $O=0.99$. For the pair of dark dotted CCPs in Figure 2 (Group Z), the value of the grouping criterion is 0.34, and the values of its component are $P=0.03$, $D=0.56$, $SL=0.10$, $ST=1$, and $O=0$.

Figure 2

5.2. Part segmentation

Now that the grouping criterion has been defined, part segmentation itself can be described.

Figure 2 illustrates the part segmentation process. Each arrow in the diagram represents a step in the process. The segmentation begins by first removing the CCPs that are outside the outline of the object (Step 2a). Then, grouping attempts that take into account the segmentation stage are

made (Step 2b). The set of possible groupings depends on whether we are at Stage 1 or Stage 2 as explained below. If a valid group (see below) is found, the boundary of the part is completed (Step 2c). If the segmentation process is at Stage 1, the CCPs making out the part are removed (Step 2d). The segmentation process cycles between the three last steps (Steps 2c, 2d, and 2e) until no more groups can be found. The segmentation process is ended (Step 2f).

The next subsections describe in details CCPs grouping and group validation (Step 2b), boundary completion (Step 2c), part CCPs removal (Step 2d) and the difference between the two segmentation stages.

5.2.1. CCPs grouping

Grouping attempts are made from all the available constant-curvature primitives (CCPs). The availability of some CCPs depends on the segmentation process stage. At Stage 1, only the outline CCPs are available for grouping, whereas at Stage 2, interior CCPs are also available. Grouping starts by selecting the longest (most salient) CCP available. Before selecting the longest CCP, cocurvilignity (colinearity or cocircularity) is used to reconstruct fragmented CCPs. A reconstructed fragmented CCP or a CCP that does not need reconstruction is called a reconstructed CCP (RCCP). For each grouping attempts, the longest RCCP available (seed RCCP) is tentatively grouped with all the RCCPs available. The quality of each grouping attempt is measured by the value of the grouping criterion. The pair of RCCPs that gets the highest value for the grouping criterion, which is also higher than the grouping criterion threshold, is considered as a possible group. The grouping criterion threshold is the minimum grouping criterion value considered for a valid group.

5.2.2. Group validation

Group validation ensures that the groups made are the ones in which we have the most confidence. We have more confidence in groups built from outline RCCPs and groups that bound the object. For groups made of two outline RCCPs, the impact of the interior RCCPs on these groups must be investigated. This is the goal of Verification 1 of the group validation step.

Figure 3

Figure 3 shows the three possible situations in which a group of two outline RCCPs can be formed. Although the two RCCPs respect the grouping criterion, they do not necessarily correspond to the best group (see Figure 3, Situation A). Verification 1 determines to which situation a grouping corresponds and rejects or accepts the group accordingly. Groupings that correspond to Situations B and C (see Figure 3) are accepted. Groupings that correspond to Situation A are rejected. The grouping situation is established by reattempting to group the outline RCCPs under validation with interior RCCPs. If no group is found, the group under validation corresponds to a grouping in Situation C. If the first two RCCPs under validation are grouped with the same interior RCCPs, the group under validation corresponds to a grouping in Situation B. In all other cases, the group under validation corresponds to a grouping in Situation A. Verification 1 applies only to groups made during Stage 1.

A match may also not be valid because the pair of matched RCCPs bounds a portion of the scene background. Verification 2 of the group validation step addresses this issue. Note that since only the exterior outline of the object is known, this verification does not handle regions of background seen through the object. However, it could be applied unmodified to handle this case if internal outlines were detected.

To verify if the interior of the object is between the two matched RCCPs as it should, coordinates of points, sampled between the two RCCPs, are verified to confirm if they belong to the set of coordinates of points of the inside area. Since the outline of the object is known, the area it covers is known and so is the set of points inside this area. Figure 4 shows how the sampled points are chosen. If all sampled points belong to the set of points of the inside area, the two RCCPs bound a region of the object, and hence, the grouping is accepted under this verification. This verification applies to both stages.

Figure 4

Even though quite powerful in selecting proper groupings for object parts, these two generic contextual verifications are typically absent from previously proposed grouping methods.

If a grouping is rejected for one of the two verifications, this pair of RCCPs is removed from the solution space along with all the pairs made of the current seed RCCP. A new grouping attempt is initiated with the next longest (more salient) RCCP available.

5.2.3. Boundary completion

Once a valid group is found, the complete boundary of the part is constructed from the set of CCPs of the object or from virtual straight-line segments. This is done by attempting to make paths of CCPs that link pairs of endpoints of the two grouped RCCPs. The paths are constructed based on three criteria. First, all the CCPs of a given path must remain inside the area defined by the lines passing by the endpoints of each of the two grouped RCCPs. The dotted path in Figure 2 (Path Y) is invalid because of this criterion. Next, two consecutive CCPs in the path must be oriented as much as possible in the same way. Finally, CCPs on the outline are chosen first. The path is constructed using proximity between endpoints as explained next.

To add a CCP at a given point in the path, a set of nearby CCPs is first built. Then, the orientation criterion is applied to select the CCP to add to the path (outline CCPs first). RCCPs can be added to the path if they reach directly the destination RCCP (the RCCP to reach to complete the path), and if they bridge a gap. During the construction of the path, there is no backtracking. Hence, if the path does not reach the destination RCCP, the path construction fails. In this case, the two unlinked endpoints are linked by a virtual straight-line segment. If two virtual straight-line segments are required to link the endpoints, the correspondence is selected such that the lines do not intersect and are both as short as possible. The complete boundary obtained is considered as the boundary of the projection of a part.

It is to be noted that dynamic programming cannot be used to build the paths since specific decisions are taken each time a CCP is added, concerning the search area, RCCP gap bridging and 3D interpretation of occlusions.

5.2.4. Part CCPs removal

For parts obtained at Stage 1, the CCPs that are inside the area of the part, and the RCCPs or CCPs that define its boundary are removed from the solution space (i.e. they are made unavailable for the remaining of the segmentation process).

5.2.5. Analysis of the two segmentation stages

The two segmentation stages differ mainly by the set of available RCCPs for grouping. At stage 1, only the outline RCCPs are available. For stage 2, all the RCCPs are available, but groupings are attempted first with outline RCCPs. For the second stage, the outline RCCPs are still considered more likely to be bounding a part. Furthermore, for Stage 2, only Verification 2 of the group validation step is made because no structural information is available to perform

Verification 1. In addition, at Stage 2, the part CCPs are not removed. This is because no structural information is available and hence, the groups are less certain. We rather keep all the ungrouped RCCPs, consider all the possibilities, and obtain overlapping parts, than make erroneous premature decisions. These overlapping parts will have to be further processed to determine the best ones. Note, however, that the seed RCCP is made unavailable in both stages to avoid redundancy. The segmentation proceeds to Stage 2 when no more groups can be formed at Stage 1.

6. Results and analysis

To validate our approach, three experiments have been designed, each investigating a specific aspect. The first experiment verifies the validity of our grouping criterion by processing synthetic constant-curvature primitives (CCPs) drawings. The second experiment extends the first experiment, by verifying the validity of our grouping criterion for images of real scenes. The third experiment investigates the consistency of the segmentation for several images of the same object. The next subsections present each of these experiments, followed by a discussion about the parameters involved in the segmentation process and computation times.

Figure 5

6.1. Experiment on synthetic CCPs drawings

The CCPs drawings chosen to validate our algorithms depict objects made of simple volumetric primitives, like cylinders, prisms, cones, etc (see Figure 5). Note that these same drawings were used to demonstrate the capabilities of PARVO [13]. Ideally, the parts obtained by our method should correspond to each of the individual simple volumetric primitives composing the object if their projections are bound by at least two outline CCPs. For projections bounded by less than

two outline CCPs, the individual faces of the volumetric primitives should be found. Our experiments show that it is totally the case for twenty-one of the twenty-three CCPs drawings tested. The segmentation obtained for the drawing of a stepladder is shown in Figure 6. Parts labelled A and B correspond to the two visible faces of the projected prism. The prism is not extracted as a whole because it has no CCPs on the outline.

Figure 6

Figure 7 shows the results for the two CCPs drawings that were not segmented perfectly. For the pen, the reason is the failure to reconstruct a fragmented circular arc (indicated by two arrows in Figure 7). This reconstruction failure causes the under-segmentation of the body of the pen and the clip. Note that for this view of the pen, the clip could not have been segmented from the body even with a perfect reconstruction since it is almost completely inside the body. The area of the clip inside the body is considered as texture by our method. The remaining CCPs of the clip would not form any group, and hence it would not be segmented. For the briefcase, the handle, which is made of volumetric primitives that are not bounded by the outline, is not properly segmented because of the grouping criterion. To obtain a good result for this CCPs drawing, the similarity in length and the distance between the CCPs should not have the same weight in the grouping criterion for different regions of the handle. Since a unique grouping criterion is used for the entire image, it is not possible to perfectly segment the handle. This problem occurs because the handle is made of many small CCPs of about the same length that are all nearby. Adapting the grouping criterion to the scale of the CCPs could alleviate the situation.

Figure 7

This experiment demonstrates that our segmentation method can process synthetic CCPs drawing exactly as expected for over 90% of the drawings we have tested. For the drawings that were not entirely processed correctly, the problems causing the errors are well understood, and affect only local portions of the objects. Therefore, our method is suited to segment objects and can be further validated with images of real scenes to observe its performance in presence of noise, texture, and shadows.

6.2. Experiment using images of real scene

The goal of this experiment is to show that our grouping criterion and our method can be used to process images of complex objects in real scenes. For this purpose, twenty-four images of objects in real scenes have been processed. Among these, nineteen were segmented properly. Segmented properly or correctly means the following. The parts obtained cover well the projection of the object and there are few overlaps that do not result from occlusions. Each part obtained can be interpreted as simple volumetric primitives and is coherent with the viewpoint. Finally, the parts obtained seem to correspond in majority to the part enumerated by a human to describe the object.

Figure 8

Figure 8 shows the part segmentation result obtained for the image of a compass. The partial ring at the top of the compass is extracted as a disc, because circular arcs describing the interior boundary of the ring are missing. The method has performed adequately with the CCPs available in this region.

Although the results obtained are not perfect for all the images tested, the majority of the parts obtained correspond to simple volumetric primitives making out the object. Hence, in general, the method reaches its goal to segment an object into its simple volumetric primitives. It succeeds for about 80% of the images. For the images that are not segmented correctly, the errors are only local to a portion of the object. The segmentation method handles relatively well textures and noise, as fragmented CCPs are reconstructed and the outline is used to obtain clues about the structure of the object. The following paragraphs discuss some typical segmentation errors.

Figure 9

Figure 9 shows the result for an airplane with a few segmentation errors. Parts corresponding to wings (Parts A, C, D, F, G, H, J, K) and an engine (Part I) have been segmented correctly. The fuselage is over-segmented (parts B and E). This is because the reconstruction of the fragmented boundary of the fuselage has failed. There is a spurious background part (part L). This is due to an error produced by the object detection step. Therefore, the presence of this part is not a failure of the grouping criterion as it has processed it as if it was a part of the airplane, which could indeed be possible. Note, however, that this error has not affected the segmentation of other regions of the airplane. Our grouping criterion can consequently tolerate object detection errors, as they remain in general local. Finally, there are parts (parts M, N, O) that correspond to markings on the airplane. These parts are found because they are in-between two simple volumetric primitives of the airplane. In this case, our grouping criterion behaved correctly in the circumstances. The CCPs composing these parts were available for grouping, and they have been grouped as it should. However, these parts are irrelevant for the modelling of the structure of this particular airplane. In this case, they could be removed by comparing their size with the sizes of

the other parts of the airplane and by analyzing their location (they are in-between two nearby more salient parts).

Figure 10

Figure 10 shows the result for another image with some segmentation errors. Aside from a leg that has been over segmented (parts E, G) because of two very close parallel straight-line segments in the area, two legs (part B) and two rungs (part F) are each extracted as one part. This is because of the viewpoint. The two individual parts become overlapped, and they appear as a prism for which two faces are visible. Since CCPs of the two individual parts are on the outline, are mutually grouped, and are validated by Verification 2, they are extracted as a whole.

Although these are not desirable parts, our segmentation algorithm is working correctly. High-level reasoning is needed to prevent this kind of grouping. This shows that our segmentation method is not completely independent of the viewpoint. However, for our image query application, a total independence from viewpoints is desired, but not essential. Obtaining consistency of description for a wider range of viewpoints than it is possible with 2D contour models was the intent for the design of this segmentation method. We believe that description by simple volumetric parts allows better viewpoint independence for comparing 2D images of objects in a database. The third experiment precisely investigates the consistency of description under changes of viewpoints.

6.3. Experiment on segmentation consistency

For modelling an object by a graph of parts, a method that can segment objects in images into parts is not enough. That method must give consistent segmentations for different images of the same object and for a certain range of viewpoints. This important issue is poorly addressed in the

previous works on part-based description of objects. Our last experiment verifies whether our segmentation method gives consistent segmentations. Fifteen images of a desk lamp have been processed. They are shown in Figure 11, along with the two types of segmentation obtained. The segmentation of the lamp pole in one or more parts is not considered as a difference in the type of segmentation since the obtained parts can easily be merged (see [2]). The significant difference in the segmentations for the fifteen images arises on the lampshade. For fourteen of the fifteen images, it is segmented as one part, and for the remaining image, it is segmented as two parts. The lampshade is segmented as one part for the majority of the images because the circular arc marked with an X in Figure 11 is the longest CCPs of the lampshade. Hence, it is grouped first and it gives a valid group with the opposite circular arc. For the image that does not give this segmentation, the circular arc marked with an X is still tentatively grouped with the opposite circular arc, but this group is not valid under Verification 1. This is because there is an additional circular arc between the two CCPs that correspond to the light bulb visible under the lampshade.

Figure 11

This experiment shows that our method is capable, in general, of producing consistent segmentations for different images of the same object. The views of the lamp used in this experiment are quite varied, and the 2D contour shapes of the lamp quite different. In some cases, texture (like the light bulb) can cause different segmentation. However, these segmentation errors remain local.

6.4. Segmentation parameters and computation times

To segment images with our method, three parameters must be set. They all involve the reconstruction of CCPs. They are the cocircularity tolerance, the colinearity tolerance, and the gap accepted between two CCPs forming a reconstructed CCP. The cocircularity and colinearity tolerances concern the perpendicular distance between two circular arcs or two straight-line segments having the same orientation. These tolerance parameters adjust the threshold in pixels of the maximum perpendicular distance allowed. The threshold on the gap is expressed as the percentage of object CCPs support for a reconstructed CCP. If these thresholds are too high, any two CCPs can form a reconstructed CCP. If the thresholds are too low, then fragmented CCPs cannot be reconstructed.

For synthetic line drawings, unique values can be used for these parameters. However, for images of real scenes, these parameters must vary from one image to another, although the range of values to obtain identical segmentations overlap for several images. This is particularly the case for the gap and colinearity parameters (Table 1). We believe that the selection of the thresholds could be automated, using constraints on the number of parts and the number of overlaps, and by analyzing the variation in the dimension of parts in specific areas when the threshold values are changed. This is a topic for future research.

Table 1

Another important issue is the computation time. In the context of a database query engine, the query image must be processed rapidly. Without optimization, the computation times obtained on an Athlon Thunderbird 1.2 GHz are all under 10 seconds (see Table 1). The complexity of the method is in the order of n^2 , where n is the number of CCPs. Note, that each time a group is

found, CCPs are removed from the solution space. A hierarchical extraction of parts using multi-scale CCP maps is another area for future research.

7. Conclusion

This paper has presented a new method to segment an object from an image of a real scene into parts that correspond to the projections of simple volumetric primitives. These projections of simple volumetric primitives can be described mainly by their two main sides. They are obtained by grouping two constant-curvature primitives (CCPs) that respect a criterion involving simple geometric relationships, like symmetry, length, distance, similarity and overlap. The outline of the object is also used to add to the grouping process knowledge about the structure of the object.

The results obtained show that our method performs well enough to be used in the context of our image database query engine. In general, the segmentations obtained for different images of the same object are consistent. That is, for the same object, the segmentation is usually the same for a range of viewpoints that are not too odd. This is in agreement with perceptual experiments that have shown that familiar views of objects are more easily segmented into parts by humans [1].

The method is capable of segmenting objects made of well-defined simple volumetric primitives, like cylinders and prisms. It can also tolerate a certain amount of deformation of the ideal shapes in addition to markings and texture on the object surface.

Work is under way to apply these results to query images in a database. First, qualitative volumetric primitives are to be inferred from extracted projections of simple volumetric primitives and a qualitative volumetric primitives graph is to be constructed. Then, a method to compare these graphs is to be implemented. Finally, the performance of the complete system will

be characterized by using the query engine on a variety of images. All these aspects are currently designed and implemented (see [2],[18]).

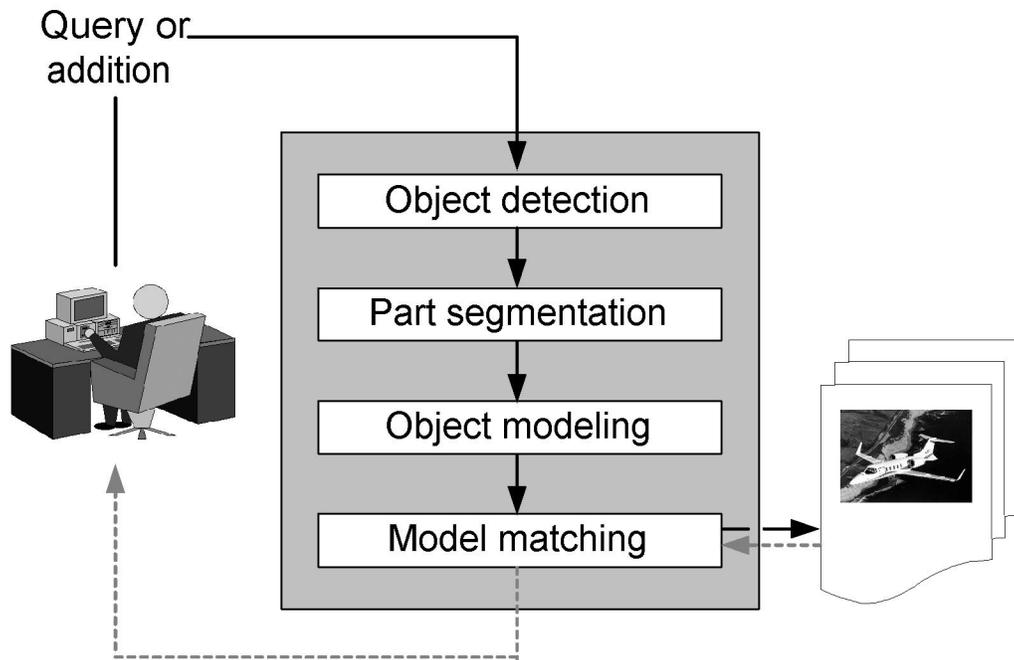
Acknowledgements

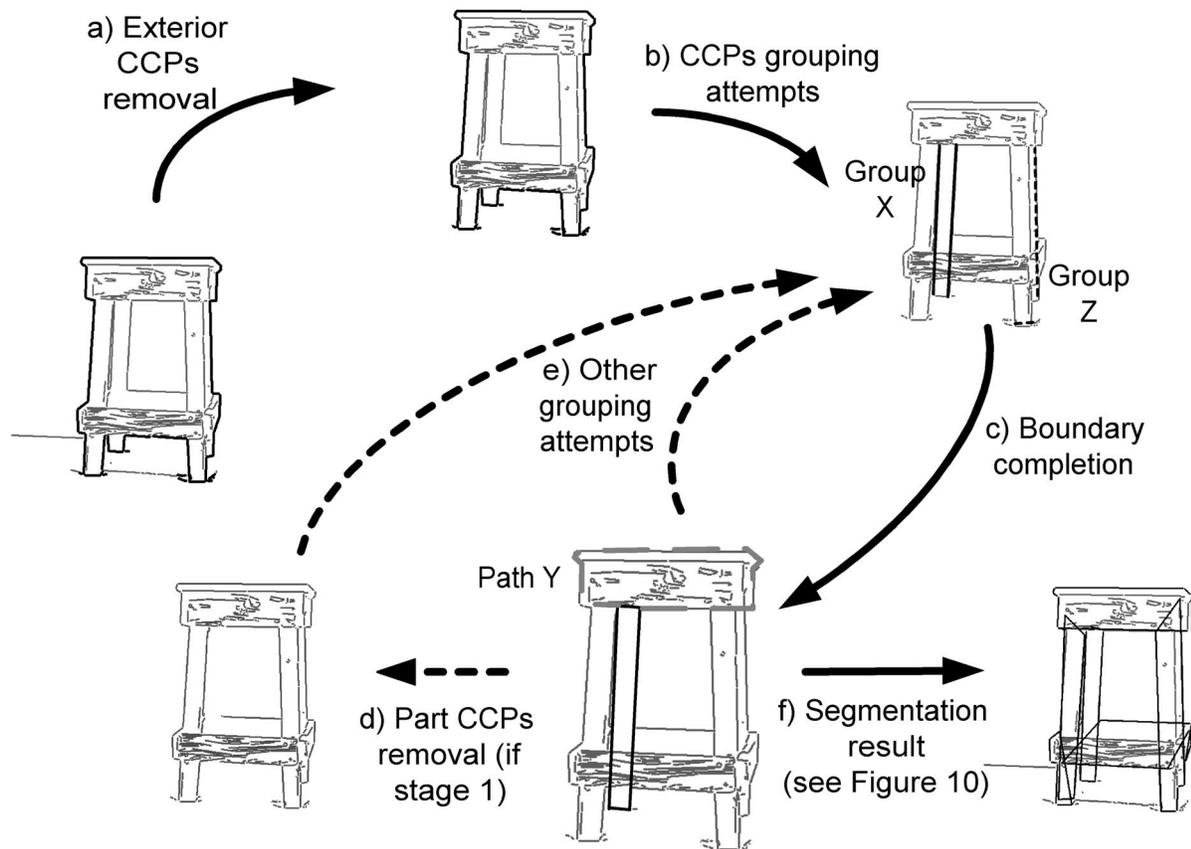
This work is supported by a postgraduate scholarship from the Natural Sciences and Engineering Research Council of Canada (NSERC) and from le Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (FCAR).

References

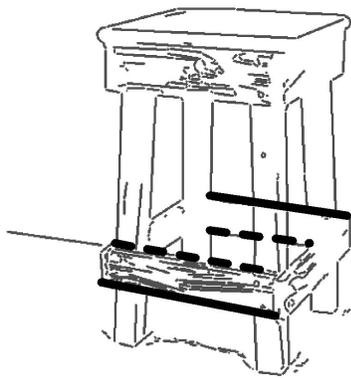
- [1] Biederman, I., Human Image Understanding: Recent Research and a Theory, *Computer Vision Graphics Image Processing*, 32 (1985), p. 29-73.
- [2] Bilodeau, G.A. and R. Bergevin, Generic Modeling of 3D Objects from Single 2D Images, in *International Conference on Pattern Recognition*, Barcelona, Spain, 2000, p. 770-773.
- [3] Bergevin, R. and M. Mokhtari, Multiscale Contour Segmentation and Approximation: An Algorithm Based on the Geometry of Regular Inscribed Polygons, *Computer Vision and Image Understanding*, 71(1) (1998), p. 55-75.
- [4] Tek, H., P.A. Stoll, and B. Kimia, Shocks From Images: Propagation of Orientation Elements, in *Computer Vision and Pattern Recognition*, Puerto Rico, 1997, p. 839-845.
- [5] Siddiqi, K. and B.J. Kimia, Parts of Visual Form: Computational Aspects, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(3) (1995), p. 239-251.
- [6] Kelly, M.F. and M.D. Levine, Region-Based Grouping Operations for Locating and Describing Objects, in *Vision Interface'95*, Québec, Canada, 1995.
- [7] Zerroug, M. and R. Nevatia, Volumetric Descriptions from a Single Intensity Image. *International Journal of Computer Vision*, 20(1/2) (1996), p. 11-42.
- [8] Bennamoun, M. and B. Boashash, A Structural-Description-Based Vision System for Automatic Object Recognition, *IEEE Transaction on Systems, Man and Cybernetics - Part B: Cybernetics*, 27(6) (1997), p. 893-906.
- [9] Pilu, M. and R.B. Fisher, Model-Driven Grouping and Recognition of Generic Object Parts from Single Images, *Robotics and Autonomous Systems*, 21 (1997), p. 107-122.
- [10] Jacot-Descombes, A. and T. Pun, Asynchronous Perceptual Grouping: From Contours to Relevant 2-D Structures, *Computer Vision and Image Understanding*, 66(1) (1997), p. 1-24.
- [11] Jacobs, D, Robust and Efficient Detection of Convex Groups, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1) (1996), p. 23-37.
- [12] Lowe, D.G., *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, 1985.

- [13] Bergevin, R. and M.D. Levine, Generic Object Recognition: Building and Matching Coarse Descriptions from Line Drawings, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(1) (1993), p. 19-36.
- [14] Dickinson, S.J., A.P. Pentland, and A. Rosenfeld, 3-D Shape Recovery Using Distributed Aspect Matching, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2) (1992), p. 174-198.
- [15] Singh, S., G.D. Seyranian, and D.D. Hoffman, Parsing Silhouettes, *Perception and Psychophysics*, (61) (1999), p. 636-660.
- [16] Bilodeau, G.A. and R. Bergevin, Outline-based Part Segmentation using Intermediate-level Symmetries, in *Vision Interface'99*, Trois-Rivières, Canada, 1999, p. 66-73.
- [17] Møktari, M. and R. Bergevin, Multiscale Compression of Planar Curves using Constant Curvature Segments, in *14th International Conference on Pattern Recognition*, Brisbane, Australia, 1998, p. 744-746.
- [18] Bilodeau, G.A. and R. Bergevin, Modeling of 2D Parts Applied to Database Query, in *Vision Interface 2001*, Ottawa, Canada, 2001, p. 228-235.

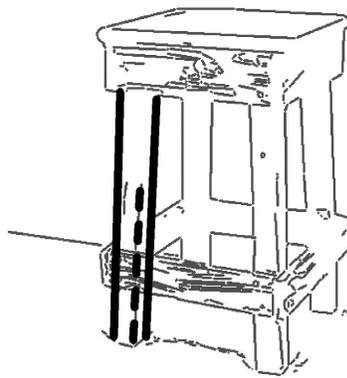




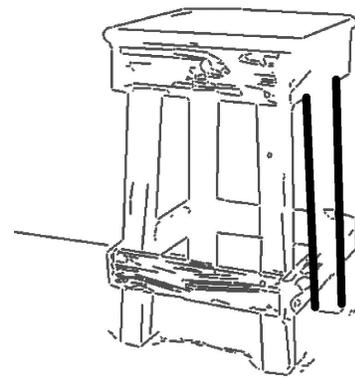
--- validation groups with interior CCPs
— initial groups



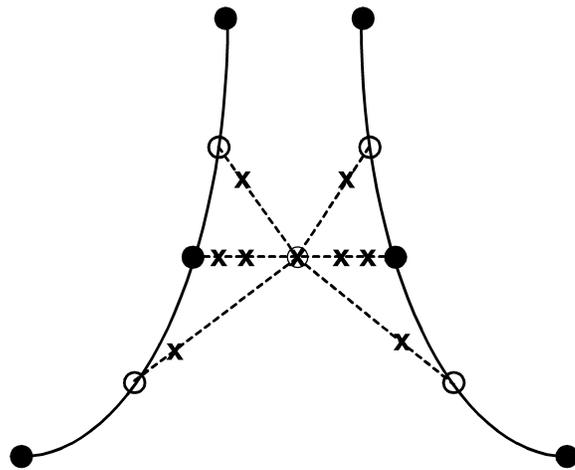
Situation A



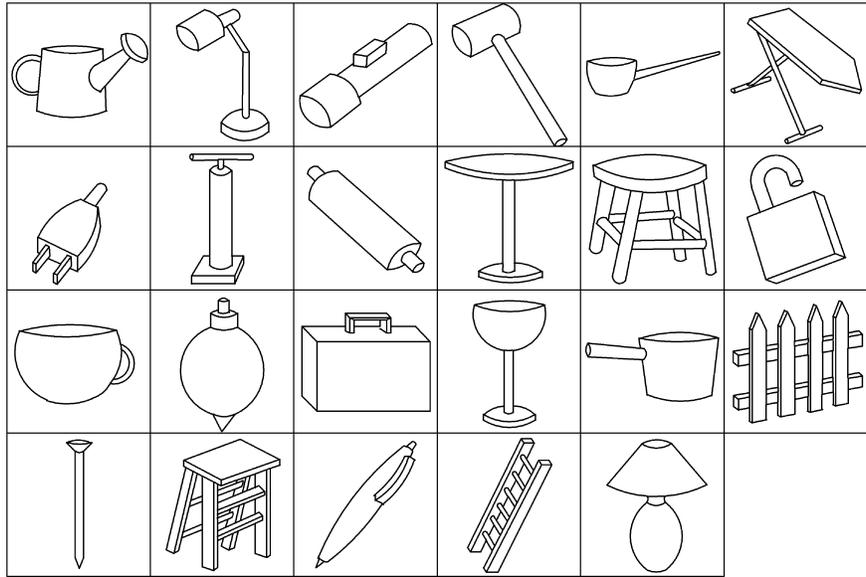
Situation B

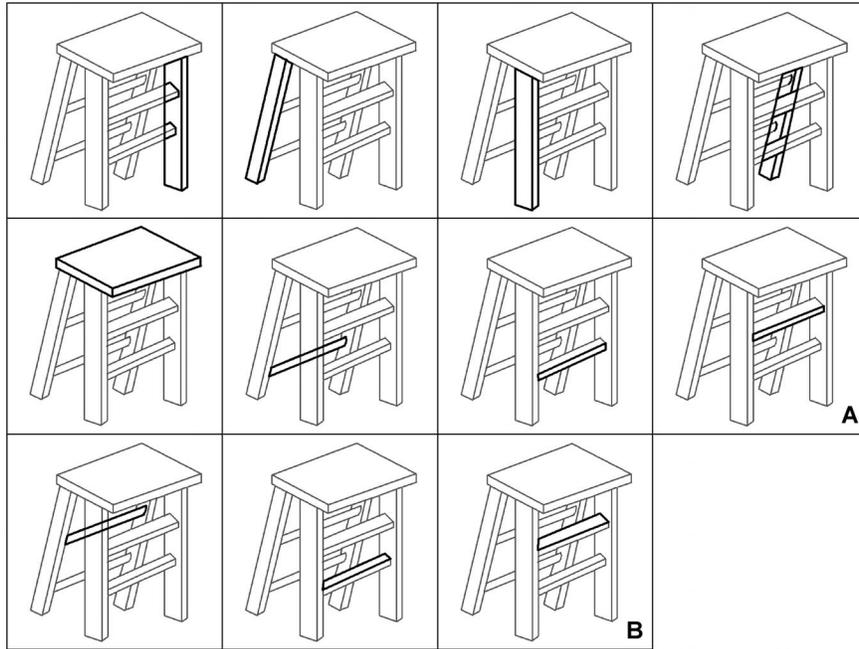


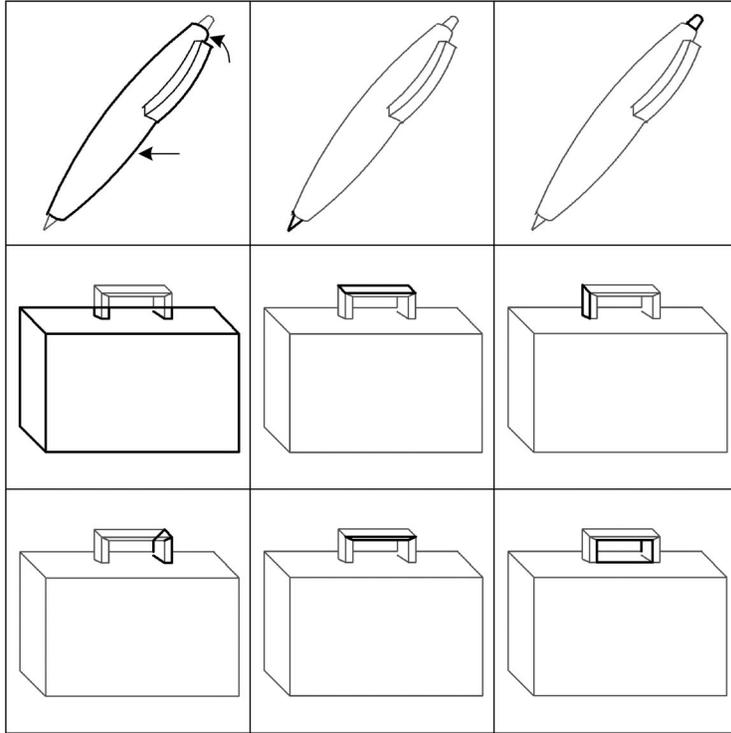
Situation C

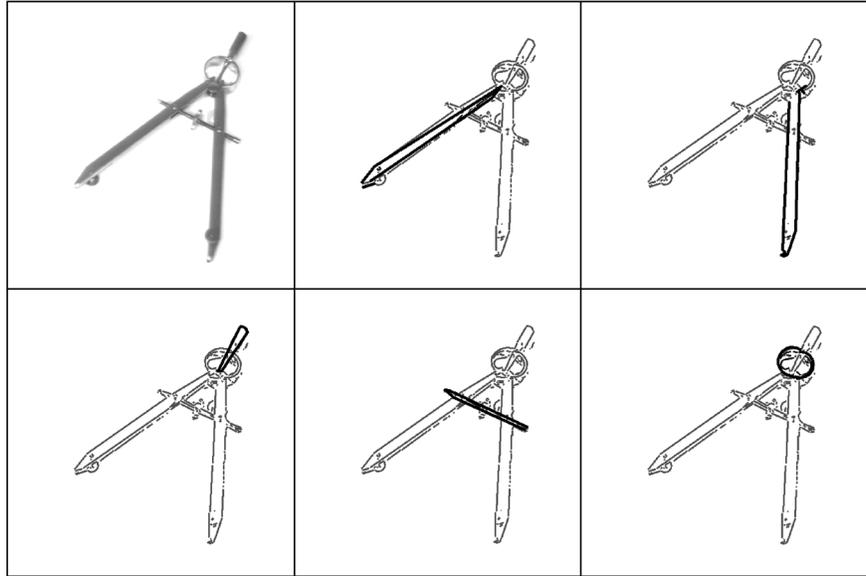


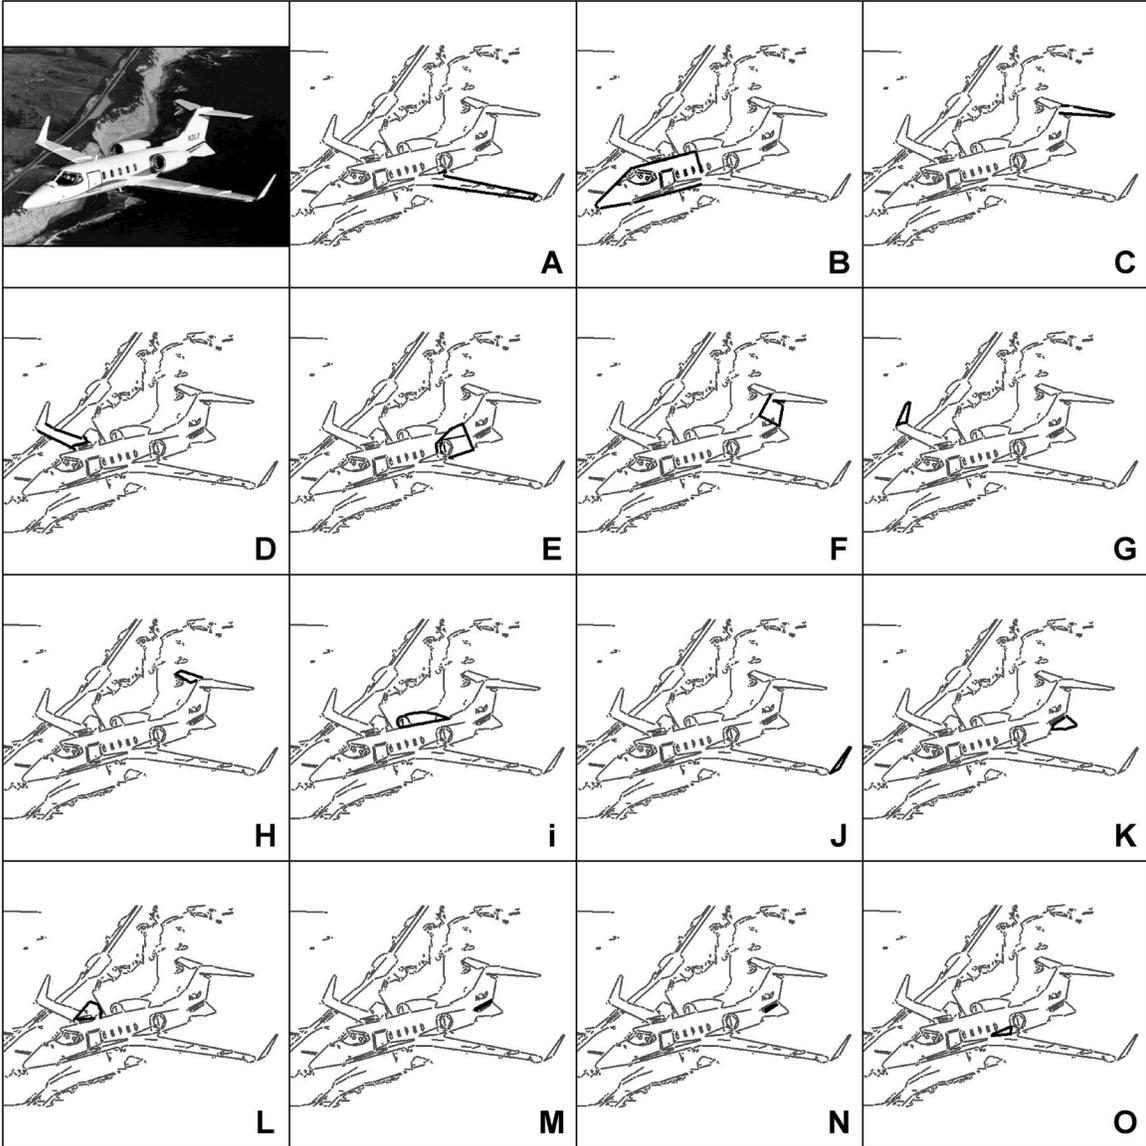
- Points describing the CCP
- Intermediate points
- x Validation points
- ⊗ Intermediate and validation point

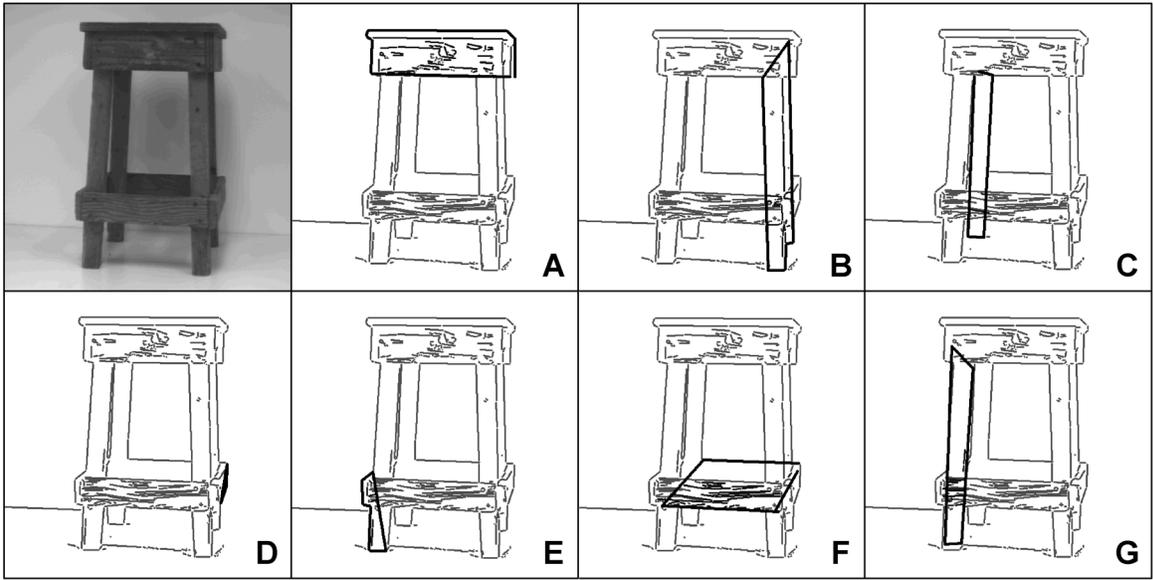












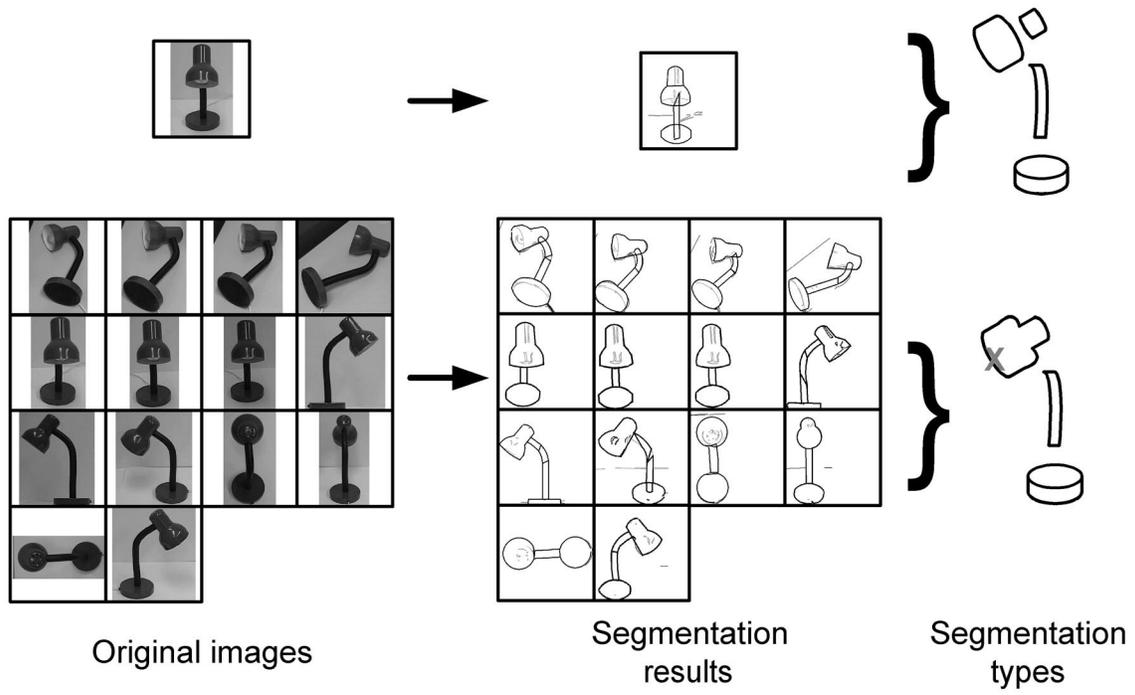


Figure 1: Overview of the query engine.

Figure 2: The segmentation process

a) The segmentation begins by removing exterior CCPs, b) CCPs grouping attempts are made, c) The boundary of the part is completed, d) The CCPs making the part are removed (if Stage 1), e) Other grouping attempts are made, f) Segmentation result.

Figure 3: Validation with interior CCPs

For Situation A, two parts with only one CCP on the outline are grouped. If both CCPs are re-grouped with interior CCPs, the group becomes invalid since they are both best grouped with other CCPs. For Situation B and Situation C, the groups are validated because both grouped CCPs are grouped with the same interior CCP or they are again grouped together.

Figure 4: Verifying if the object is between to CCPs.

Intermediate points are midpoint between CCPs points. Validation points are midpoint between intermediate points or between an intermediate point and the midpoint of a CCP.

Figure 5: The line drawings used in the first experiment (from [13]).

Figure 6: Segmentation result for a stepladder.

Figure 7: Segmentation results for two problematic drawings

Figure 8: Segmentation result for a compass image.

Figure 9: Segmentation result for an airplane image.

Figure 10: Segmentation result for a stool image.

Figure 11: Two types of segmentation for images of the same desk lamp.

Table 1: Range of parameter values to obtain identical segmentations.

Image name	Number of CCPs	Cocircularity tolerance (pixels)	Colinearity tolerance (pixels)	GAP (%)	Computation time (s)
Learjet1	171	[0,1]	3	[20,30]	7.14
Learjet2	70	[0,3]	[3,5]	[40,80]	1.64
Compass	92	4	[7,8]	[40,50]	2.90
Stool2	177	[0,4]	7	[40,50]	7.62
Cup	39	[4,5]	[1,20]	[40,50]	0.83
VR-Room_Lamp1	61	[1,3]	[1,14]	[20,50]	1.27
VR-Room_Lamp2	41	[6,11]	[1,20]	[20,30]	0.65
VR-Room_Lamp3	37	[3,6]	[1,20]	[20,40]	0.64

Summary

This paper presents one significant aspect of a project aimed at the design of an image database query engine, where the images are searched at the 3D object-level. This approach is a novelty since the majority of existing image database query engines search images by comparing the colours, the textures and the 2D shape of regions in the images. The main aspect and contribution of this paper is a new method for object part segmentation, from circular arcs and straight-line segments primitives, that is suited to process images of object in real scene. In the context of this work, 2D parts are defined as regions delimited by groups of circular arcs and straight-line segments (constant curvature primitives), which can later on be interpreted as the projections in the plan of simple volumetric primitives, like cylinders and prisms. Our new method is based on perceptual grouping. It groups constant curvature primitives (CCPs) using intermediate-level geometric relationships and object outline information. The use of the outline of the object reduces the complexity of the grouping process, since many groups are discarded based on this information. The principle of our method for part segmentation is to group two CCPs with a criterion based on intermediate-level geometric relationships and with the outline of the object to account for its structure. Subsequently, the complete boundary of the part is detected by making cycles with object CCPs to link the two grouped CCPs. The shape of the parts extracted is only limited by the structure of the object (obtained via the outline) and by the geometric relationships between the two CCPs grouped. We will show that our approach is suited to process 3D objects in 2D images of real scenes by several validating experiments.

Biographical sketch

GUILLAUME-ALEXANDRE BILODEAU received the B.Sc.A. Degree in Computer Engineering from Université Laval, Canada, in 1997, and the M.Sc.A. Degree in Electrical Engineering from the same University in 1999. He is currently a Ph.D. candidate at the Computer Vision and Systems Laboratory of Université Laval. His research interests include object recognition, image database query and graph matching. He is a member of Quebec's association of professional engineers (OIQ), and a member of CIPPRS.

ROBERT BERGEVIN received the B.Eng. degree in Electrical Engineering and the M.A.Sc. Degree in Biomedical Engineering from École Polytechnique in Montréal, in 1982 and 1984, respectively, and the Ph.D. degree in Electrical Engineering from McGill University in 1990. From May 1984 to September 1985, he was a practicing clinical engineer at Sacré-Coeur Hospital in Montréal. From March to August 1990, he worked as a researcher in the field of automated inspection at the Production Engineering Research Laboratory of Hitachi Ltd. in Yokohama, Japan. His research interests encompass computer vision and robotics with an emphasis on object and scene description, modeling and recognition. Since November 1990, he has been with the Computer Vision and Systems Laboratory at Laval University, Quebec City, where he is currently a Full Professor in the Department of Electrical and Computer Engineering. Dr Bergevin is a member of the Province of Quebec's association of professional engineers (OIQ), a member of the IEEE Computer Society, and CIPPRS, and an Associate Editor of Pattern Recognition, the Journal of the Pattern Recognition Society.