MVA2000 IAPR Workshop on Machine Vision Applications, Nov. 28–30, 2000, The University of Tokyo, Japan

3—24

Qualitative Decomposition of Range Images into Convex Parts/Objects

Xiaoyi Jiang *
Department of Computer Science
University of Bern

Abstract

Previous works on range image segmentation concentrate on surface patches that can be well represented by certain mathematical functions. In this paper we consider the more qualitative segmentation problem of decomposing a range image into convex parts/objects. An edge-based approach is proposed which uses an adaptive contour closure algorithm in conjunction with a global convexity test. Experimental results on two range image sets are reported to demonstrate the performance of our convex decomposition technique.

1 Introduction

Typically, range images are segmented into surface patches that can be well represented by certain mathematical functions, e.g. bivariable polynomial functions, quadrics, and superquadrics. In this paper we consider the more qualitative segmentation problem of decomposing a range image into convex parts/objects.

Convexity is a powerful grouping property for several reasons. It is unlikely that a random set of data will form a convex shape. In addition grouping based on convexity can be done independent from any mathematical shape model in a qualitative way. This potentially enlarges the applicability of such an approach compared to model-based segmentation.

In this work we assume that objects are constituted by convex parts and these parts are separated from each other by concave contours. An example of such an object is given in Figure 1 (left) which consists of two boxes. On the other hand Figure 1 (right) shows an L-shaped object that violates our assumption. This object can be considered as two boxes meeting at their contours. Despite of the assumption the class of objects under consideration in this work is general enough to be useful in various applications, particularly when dealing with objects like boxes and drums [5], and mail pieces [1, 6]. Interestingly, it turns out that the decomposition algorithm proposed in this paper is able to process

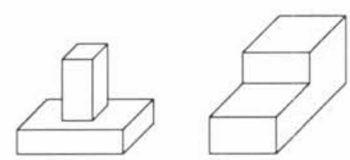

Figure 1: Object consisting of two boxes (left) and L-shaped object (right).

even objects beyond the defined class, for instance, objects like the L-shaped one in Figure 1, properly.

In the literature very few work has been reported on convex decomposition of range images. In [8] simulated electrical charge distributions over an object's surface are used for this purpose. Trucco [7] deals with slice data and partitions each slice into convex segments. Then, the segments are locally grouped into larger parts. Since no global convexity test is carried out, there is a high risk of obtaining non-convex parts.

In contrast we consider a global convexity test as an essential component of convex decomposition. Our approach is edge-based and uses an adaptive contour closure algorithm within the hypothesize-and-verify paradigm, where the verification step is totally guided by a global convexity test. We describe the adaptive contour closure algorithm and various components of our convex decomposition approach in Sections 2–4. Then, experimental results are reported in Section 5. Finally, some discussions conclude the paper.

2 Contour closure algorithm

Fundamental to convex decomposition is a convexity test that decides if an image region represents a convex part. Given such a convex test, we may apply traditional region-based approaches such as region growing to partition a range image into convex parts. In this work, however, we prefer an edge-based approach due to its various advantages [3].

For our purpose we need to detect edge points on concave contours. The basic difficulty we encounter then in solving the convex decomposition problem

*Address: Neubrückstrasse 10, CH-3012 Bern, Switzerland. E-mail: jiang@iam.unibe.ch

lies in the incomplete nature of detected contours. Generally, edge detection methods cannot guarantee closure of boundaries for surface extraction, thus resulting in the need of a subsequent grouping and completion process.

In [3] we have developed an adaptive contour closure algorithm for this purpose. The fundamental observation there is that any contour gap can be closed by dilating the edge map, or equivalently eroding the regions. If the largest gap of a region has a length L, then $L/2$ erosions will successfully complete the region. However, one has no idea about the actual value of L before the grouping process is finished. In order not to miss any region, we potentially have to select a high value for L as maximum allowable gap length, resulting in a consistently large number of erosions applied to all regions of an input image. This is not only an unnecessary overhead in dealing with regions that are (almost) closed. But also relatively small-sized or thin regions will disappear. Instead, the adaptive contour closure algorithm carries out only the minimum number of erosions necessary for each particular region. Our method (see Figure 2 for an outline) is embedded in the hypothesize-and-verify paradigm. It increases the number of erosions only for those regions that cannot be successfully verified.

From the input edge map, region hypotheses can be found by a component labeling. Usually, this initial region map contains many instances of under-segmentation. To distinguish between the correctly segmented and under-segmented regions, we perform a region homogeneity test for each region R of the initial segmentation. If the region homogeneity test is successful, the region R is recorded. Otherwise, there still exist open contours within R. In this case we perform one erosion operation within R, potentially closing the gaps in the contours. Again, a component labeling is done for R to find new region hypotheses, and these are verified in the same manner as for the initial regions. This process of hypotheses generation (component labeling) and verification (region homogeneity test) is repeated until the generated region hypotheses have been successfully verified or they are not further considered because of a region size smaller than a preset threshold T_{size}. The task considered in [3] is that of segmenting a range image into surface patches. Accordingly, the region homogeneity test is conducted using the fit error of regions by means of biquartic surface functions.

The adaptive contour closure algorithm is of general nature and can also be applied to other tasks by replacing the region homogeneity test. For the purpose of convex decomposition here this is done by a global convexity test.

```
/* Hypotheses generation */
perform component labeling on input edge map;
List = { connected regions of size > T_size };
while (List != ∅) {
  select arbitrarily a region R from List;
  /* Hypotheses verification */
  verify R using region homogeneity test;
  if (successful)
    record region R;
  else {
    /* Hypotheses generation */
    perform one erosion step within R;
    perform component labeling within R;
    List += { regions of size > T_size within R};
  }
}
postprocessing;
```

Figure 2: Adaptive contour closure algorithm.

3 Convexity test

A convexity test may be conducted by computing the 3D convex hull of the points of a given image region and then the maximal distance of the points to the convex hull. The points are regarded to form a convex part only if this maximal distance is small enough. Computation of convex hulls and distances in 3D tends to be computationally expensive. Therefore, we resort to another solution that reduces the initial 3D problem to 2D tests.

We assume that each image row/column corresponds to a curve in a plane in 3D, resulting from the intersection of the plane with the surfaces of objects in the scene. This condition is satisfied by a wide range of scanners. Among them, some scanners provide range images $z(x, y)$ regularly sampled in both coordinate directions that are particularly easy to deal with. For description clarity we will use this type of range images to introduce the convexity test. But it is easy to see that the discussion applies to other types of scanners fulfilling the condition above as well.

Given an image region R that corresponds to a convex part, each image row in R with y being a constant y_0 is simply a convex 2D curve in the xz-plane (with respect to the positive z-axis). Similarly, each image column in R with x being a constant x_0 implies a convex 2D curve in the yz-plane. As a convexity test we may verify the convexity of all rows and columns in a given image region R. In general, however, this is only a necessary but not sufficient condition for R being a convex part in 3D. A counterexample is the function $z = xy$ with negative Gaussian curvature. Here both $z = y_0x$ for a constant y_0 and $z = x_0y$ for a constant x_0 are straight lines and thus convex. But the global shape is not convex. Generally, the 2D-based convexity test will fail in the case of hyperbolic surfaces. However, it suffices for objects like boxes and drums [5], mail pieces [1, 6]. In addition, spheres, cylinders, and planes that are frequently used in manufacturing can be handled well. In our experiments reported in Section 5 we typically deal with parts that

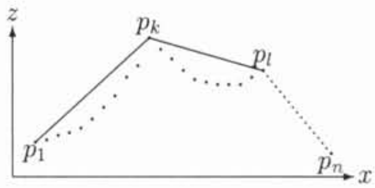

Figure 3: Convexity test of a curve by means of computation of convex hull and distances.

have non-negative Gaussian curvatures. Therefore, we can still use this simple convexity test.

It remains to solve the problem of verifying the convexity of a 2D curve. For this purpose the idea of a universal convexity test (computation of convex hull and distances) is applied to this 2D problem instance. We adapt the well-known Jarvis's march algorithm for convex hull construction of 2D points in the following way. Given n points $p_1, p_2, \ldots, p_n$ of a curve (see Figure 3), p_1 is definitely a vertex of the convex hull. The next vertex is the point p_k that provides the smallest angle between $\overline{p_1p_k}$ and the z-axis. Then, the next vertex of the convex hull is $p_l, l > k$, that provides the smallest angle between $\overline{p_1p_k}$ and $\overline{p_kp_l}$, and so on. The search is continued until we reach p_n. For each pair of two successive vertices p_kp_l of the convex hull we compute the deviation of each point $p_i, k+1 \leq i \leq l-1$, between p_k and p_l from the convex hull by means of the perpendicular distance of p_i to the line $\overline{p_kp_l}$. This way we finally obtain the maximal derivation of all points from the convex hull. For a given image region R the global maximum from all rows and columns in R is tested against a tolerance threshold for the convexity of R.

4 Postprocessing

When applying the adaptive contour closure algorithm outlined in Figure 2 to the convex decomposition problem, it remains to specify the postprocessing step. In the results we obtain without postprocessing edge pixels are not considered to be part of regions. In addition, the erosion necessary for contour closure discard pixels near region boundaries. These unlabeled pixels should be added to their corresponding regions.

As a simple strategy for doing this we can assign an unlabeled pixel p the label of an adjacent labeled pixel q if the depth (z) difference between p and q is lower than a preset tolerance. The assignment is repeated until no further operation is possible.

More rigorously, we can dilate a region once where only adjacent unlabeled pixels are considered. Then, the resulting region undergoes the convexity test described in the last section. In case of success the dilated unlabeled pixels are merged to the region. This procedure can be applied to all regions and repeated to perform all possible merge operations. Obviously, this second postprocessing method is computationally more expensive than the first one in general.

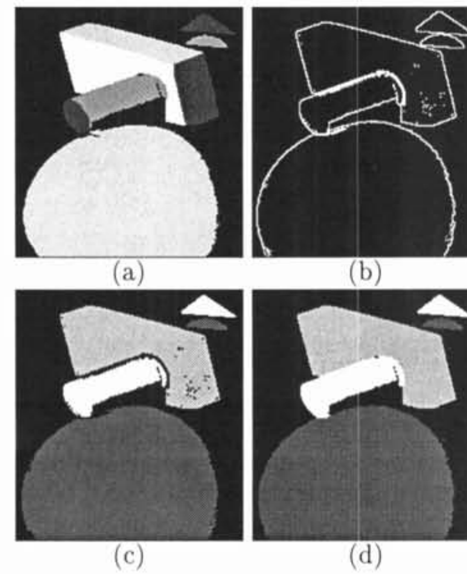

Figure 4: Image col2+ball from MSU set: (a) surface segmentation; (b) input edge map; (c) convex decomposition before postprocessing; (d) final result.

Interestingly, we have experienced that there is no essential difference in the results achieved by the two methods.

5 Experimental results

The convex decomposition algorithm has been implemented in C and tested on two range image sets. In our experiments we have used the edge detection method reported in [2] that assigns each pixel a jump edge strength value and a crease edge strength value together with the corresponding edge type (convex/concave). Only jump edge points and concave crease edge points with sufficient high edge strength values are taken to form a binary edge map as the input to the convex decomposition algorithm.

The first source of test images was the popular MSU range image set acquired by a Technical Arts scanner. The results of one test image from this set are shown in Figure 4. For comparison purpose the surface segmentation using the method described in [3] is given there as well. As expected the three surfaces of the box and the two surfaces of the cylinder are grouped together as a convex part, respectively.

The second test image set was acquired by a K2T model GRF-2 structured light scanner and has served for comparing range image segmentation algorithms [4]. The final results of four images from this set are shown in Figure 5. Here surface types cone, cylinder, and sphere are present besides planar surface patches.

As stated in the introduction section, the basic assumption we made in this work is that objects are constituted by convex parts and these parts are separated from each other by concave contours. Objects such as the L-shaped one in Figure 1 violate this assumption. A similar object can be seen in the test

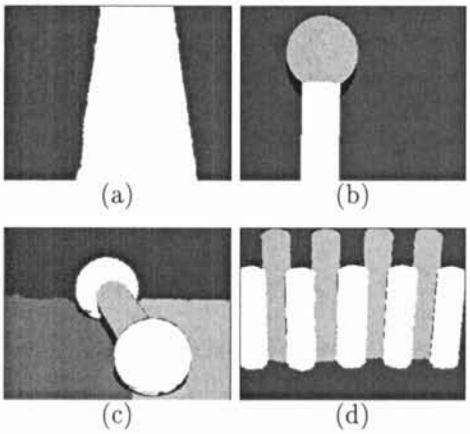

Figure 5: Images from K2T set: (a) conepb60: cone+planes; (b) balljoint: cylinder+sphere+plane; (c) dogbone2: cylinder+spheres+planes; (d) cylsocpb: cylinders+planes.

image from the MSU set given in Figure 6. In this case the separation into two convex parts partially occurs within a surface (the brightest one), resulting in a long contour gap in the edge map. The adaptive contour closure algorithm in conjunction with the global convexity test generates two relatively small convex parts which are then expanded by the postprocessing step. This example demonstrates the ability to deal with objects that violate the basic assumption of our convex decomposition method.

The computation time on a SUN Ultra5 workstation is listed in Table 1. The resolution of test images from the MSU set varies around 200×200 pixels; that of col2+ball amounts to 217×197 pixels. On the other hand, all K2T images have a uniform resolution of 480×640 pixels. The listed computation time is more or less typical for all images of each test set. As expected the rigorous postprocessing method is usually more expensive than the simple one. Since they produce essentially the same final results, we generally prefer the fast postprocessing based on simple depth value comparison.

6 Conclusions

Previous works on range image segmentation concentrate on surface patches that can be well represented by certain mathematical functions. In this paper we consider the more qualitative segmentation problem of decomposing a range image into convex parts/objects. An edge-based approach has been proposed which uses an adaptive contour closure algorithm in conjunction with a global convexity test. Experiments have been conducted on two range image sets and shown good results.

One limitation of the current implementation is the 2D-based global convexity test described in Section 3. However, it is important to emphasize that this simple test suffices for objects like boxes and drums [5], mail pieces [1, 6], and those parts frequently used in manufacturing, and is therefore useful in many practical situations. Moreover, the algorithm is general enough to be extended to other domains by simply replacing the convexity test used here through a more sophisticated one.

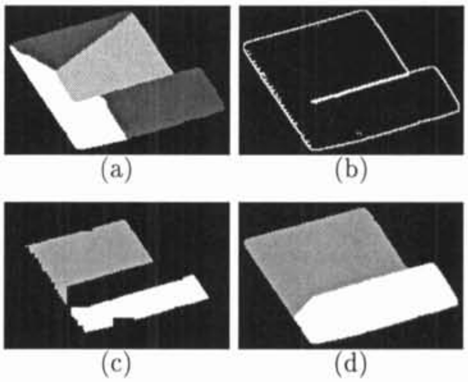

Figure 6: Image grnblk3-1 from MSU set with an object violating basic assumption: (a) surface segmentation; (b) input edge map; (c) convex decomposition before postprocessing; (d) final result.

	edge detection	convex decomp. (1)	convex decomp. (2)
col2+ball	0.2	0.2	0.8
conepb60	2.1	4.1	6.7

Table 1: Computation time in seconds: the numbers in brackets refer to the postprocessing method.

References

[1] A. Gupta, G. Funka-Lea, K. Wohn, Segmentation, modeling and classification of the compact objects in a pile, Proc. of SPIE Vol. 1192, Intelligent Robots and Computer Vision VIII: Algorithms and Techniques, 98–108, 1989.

[2] X. Jiang, H. Bunke, Edge detection in range images based on scan line approximation, CVIU, 73(2): 183–199, 1999.

[3] X. Jiang, An adaptive contour closure algorithm and its experimental evaluation, IEEE Trans. on PAMI. (accepted for publication)

[4] M.W. Powell, K.W. Bowyer, X. Jiang, H. Bunke, Comparing curved-surface range image segmenters, Proc. of the 6th ICCV, 286–291, 1998.

[5] J.C. Sluder, M.A. Abidi, Parallel range data processing: A real case study, Proc. of SPIE Vol. 1608, Intelligent Robots and Computer Vision X, 248–261, 1991.

[6] F. Solina, R. Bajcsy, Range image interpretation of mail pieces with superquadrics, Proc. of AAAI, 733–737, 1987.

[7] E. Trucco, Inferring convex subparts from slice data, PRL, 12: 707–715, 1991.

[8] K. Wu, M.D. Levine, 3D part segmentation using simulated electrical charge distributions, IEEE Trans. on PAMI, 19(11): 1223–1235, 1997.