Recovering CAD Models from Scanned Data

Tomáš Pajdla, Václav Hlaváč, Daniel Večerka
{pajdla,hlavac,vecerka}@fel.cvut.cz

Center for Machine Perception, Faculty of Electrical Engineering
Czech Technical University,
121 35 Prague 2, Karlovo náměstí. 13, Czech Republic

Abstract. The sequence of steps that allow to recover volumetric 3D CAD models of a 3D object from a set of scanned data (range images) is described. The novelty of the approach is twofolds: (i) the whole chain of providing usable technology for industry was closed, (ii) segmentation of the triangulated surface into analytic quadric patches is used to considerably improve the obtained 3D models of a typical industrial objects. The range finder is used to capture range images from different viewpoints. The triangulated surface is constructed over each range image. The data sets are reduced by decimation of triangular meshes. Surfaces are registered into a common object centered coordinate system and data are fused into a tetrahedralization of the points by α-shape. The tetrahedralization is segmented into explicitly model paths. The beautifying step improves fused data by moving the projected points onto the recovered models of surfaces. Such data enable direct reconstruction a full 3D model of the object.

1 Introduction

Rapid prototyping for mechanical engineering and other industries needs a technology that allows to create a 3D computerized geometric model from an existing 3D object. Let us mention the automotive industry as an motivating example. A designer first creates a clay model of a car body. Then, 3D coordinates of the points on the clay model surface are typically measured by a precision 3D coordinate measurement machine. Such data are the input into the further computer-aided design processing.

This paper shows a computer vision based technology that allows to automatize the process of constructing the model of a 3D free form object from a set of range images. A *range image* represents distance measurements from an observer to an object; it yields a partial 3D description of the surface from one view only. It may be visualized as a relief made by a sculptor.

Several range images are needed to capture the whole surface of an object. Each image yields a point cloud in the co-ordinates related to the range sensor, and successive images are taken in such a way that neighboring views slightly overlap, to provide information for later fusion of partial range measurements into one global, object-centered, co-ordinate system.

Range image registration finds a rigid geometric transformation between two range images of the same object captured from two different viewpoints. The recovery can either be based on explicit knowledge of sensor positions, e.g. if it is held in a precise robot arm or on geometric features measured from the overlapped parts of range data. Typically, both sources of information are used; an initial estimate of the appropriate geometric transformation can be provided by image feature correspondence, range image sensor data, an object manipulation device or in many cases by a human operator.

A *fusion of partial surface descriptions* into global object-centered coordinates requires known geometric transformations between object and sensor The process of the fusion then depends on the data representing one view, e.g. from simple point clouds, triangulated surfaces, to parametric models as quadric patches. The result of the fusion is consistent geometric tetrahedralization of the 3D points which can be afterwards converted to a triangulation.

This SOFSEM'98 short paper submission intends to point the conference audience to advanced computer vision technology. The paper sketches the whole process only briefly and describes more in detail only steps that are novel and were developed by us.

2 Model-building, an overview

This 3D model reconstruction task has been approached by several research groups in recent years, and many partial solutions were proposed, e.g. [1–3]. We will present here one of the possible approaches to the task. The method automates the construction of a 3D model of a 3D free form object from a set of range images as follows:

1. The object is placed on a turn table and a *set of range images from different viewpoints are measured* by a structured light (laser plane) range finder.
2. A triangulated surface is constructed over the range images.
3. Large data sets are reduced by *decimation* of triangular meshes in each view.
4. Surfaces are *registered* into a common object centered co-ordinate system and outliers in measurements are removed.
5. The transition from a measured point cloud to a object surface is done using α-shapes.
6. The surface is segmented into quadric patches (such shapes appear very often in the industry).
7. Knowledge yielded from segmentation allows to improve (beautify) the point clouds.
8. A full 3D model of the object is reconstructed by a surface fusion process (this step is not treated in this contribution).

Our novel contributions to the formulated problem are:

1. A practical technology that allows to solve the whole 3D reconstruction task is presented. The method is probably closer to industrial applications than most methods of others.

2. An improvement of the Iterative Closest Point (ICP) algorithm [4] for range images registration called "Iterative Closest Reciprocal Point Algorithm" (ICRP) is proposed. ICRP algorithm allows to correctly register the surfaces with large occlusions.
3. The step from point clouds to a surface that utilizes results of range image segmentation to planar, quadric or superquadric surface patches. This allows to improve (beautify) the model using the segmentation considerably.

3 Steps of model construction

3.1 Range image acquisition

Traditional approach to 3D data acquisition in mechanical engineering is by using the touching probe. The touching probe is a mechanical device allowing to measure the coordinates of a contact point between the probe and the surface. Though the touch probes are precise the mensuration by them is very tedious and slow.

Much faster acquisition can be achieved by using the Laser Plane Range Finder [5] that builds on the optical triangulation principle. A light plane is projected on the surface and its intersection with the object is observed by a camera. 3D coordinates of an illuminated points are computed by intersecting the light plane with the ray which projects the illuminated point into the camera. Each image provides one intersection curve, i.e. one planar profile is reconstructed. The whole surface can be then obtained by scanning a sequence of the profiles by moving the camera-projector rig along the object.

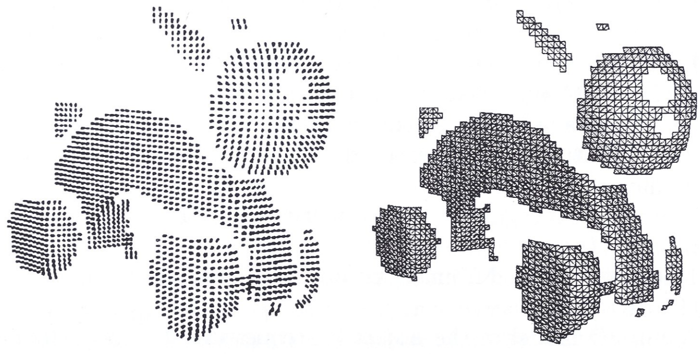

Fig. 1. (a) Points measured by the rangefinder, (b) Triangulation obtained from the four-connected mesh constructed over the points independently in each view.

3.2 Mesh construction from isolated measured points

Laser Plane Range Finder measures isolated points on a surface, Fig. 1(a), but our ultimate goal is to measure a surface passing through them. We need therefore to interpolate the points to form a surface. Very straightforward and simple method might use proximity of points to construct their triangulation. In other words, measured points are used as vertices and close points are connected so that they altogether form a triangulated mesh.

In the case of the Laser Plane Range Finder [5], a surface is reconstructed in planar scanlines. The scanlines are put one next the other by translating the camera–projector rig. This suggests to parametrize the surface across the scanlines by the motion the of camera–projector rig. The second parameter can be introduced by parameterizing the surface along the scanlines. It can be shown [5] that under some reasonable setup of the camera-projector rig, the above parameterization allows easy construction of 4–connected mesh following parametric curves just by connecting the points found in the neighboring rows of the image and the points in neighboring scans with the same image row co-ordinate. The triangulation of the mesh can then be obtained by splitting each quadrangle into two triangles, Fig. 1(b).

3.3 Decimation of triangulated surfaces

Often, we wish to reduce the number of triangles representing the visual surface in areas where its curvature is low. The data reduction is particularly useful for the registration of neighbouring views since it has in worst $O(N^2)$ complexity in number of points.

We formulate this task as looking for the best approximation of a triangulated surface by another triangulated surface that passes through a subset of the vertices of the original mesh [1]. For instance, we might look for the closest triangulated surface with maximally n triangles, or we might want to simultaneously minimize n and a residual error to get a consensus between the precision and space costs using Minimum Description Length principle [6]. Figure 2(a) shows the decimated triangulation.

3.4 ICRP registration

The *range image registration* finds a rigid geometric transformation between two range images of the same object captured from two different viewpoints. The recovery can be based (a) on explicit knowledge of sensor positions, e.g. if it is hold in robot arm or (b) on geometric features measured on the object. Typically, both sources of information are used. Initial estimate of geometric transformation can be provided by image feature correspondence, range image sensor, object manipulation device and in many cases by a human operator.

We have have adopted two approaches to solve the surface registration. The first approach uses an interaction with the user. The mutual position of two surfaces is defined by aligning three pairs of matching points. We let the user

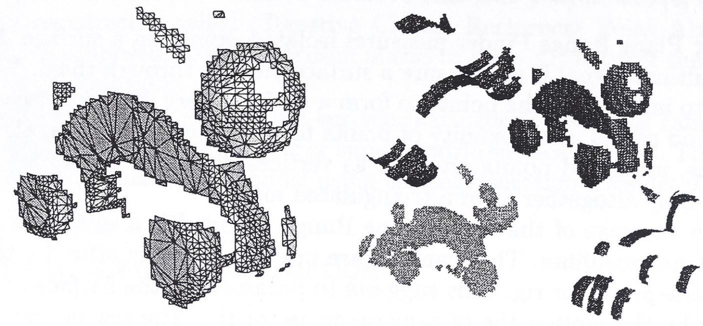

Fig. 2. (a) Decimated mesh, (b) Set of the meshes measured from different viewpoints.

to select a few point pairs (minimum is three) on the surfaces. The approximate registration is obtained by moving one of the surfaces so that the sum of squared distances between the matching point pairs is minimal. The second approach extracts special points and curves of interest from surfaces using its differential structure [7]. Such approach allows to finding an initial transformation automatically by trying to register only the points of interest. Using the points and curves of interests instead of the whole surfaces leads to considerable speeding up the registration of the surfaces without a human interaction. However, an automatic method can only be used if the surface has quite rich structure (cavities, changing curvature, etc.). That is why an interactive way of defining the initial registration is still of a practical value for simple surfaces.

The precise alignment of the data can be done automatically by a gradient minimization provided that a good starting transformations are available. The iterative closest point algorithm (ICP) developed by Besl and McKay [4] solves registration provided that a good initial estimate of transformation T is available. The ICP algorithm assumes that one of the surfaces is subset of the second. It means, that only one surface can contain points without correspondence to the second surface. We modified the ICP algorithm. Our Iterative Closest Reciprocal Point (ICRP) algorithm is able register partial corresponding surfaces. We used the method of reciprocal points [8] to eliminate the points without correspondence.

Surface registration looks for the best transformation T that overlays P and X. In other words T is found by minimization of

$$e = min_T \ \rho(P, T(X)),$$

where ρ is a function evaluating the quality of the overlap. In Euclidean geometry it might be a distance between the points on a surface.

Let us assume point $\boldsymbol{p}$ on the surface P and its closest point $\boldsymbol{y}$ on the surface X. The closest point on the surface P, to the point $\boldsymbol{y}$ is the point $\boldsymbol{r}$. The points

p, complying the condition that the distance is lower then ϵ, are ϵ-reciprocal. Only these points are registered. If set of ϵ-reciprocal points on the surface P is marked P_ϵ, that we can write the ICRP algorithm like this:

1. Initialize $k = 0$ and $P_0 = P$.
2. Find closest points Y_k for P_k and X .
3. Find reciprocal points $P_{\epsilon 0}$ and $Y_{\epsilon k}$.
4. Compute the mean square distance d_k between $P_{\epsilon k}$ and $Y_{\epsilon k}$.
5. Compute the transformation T between $Y_{\epsilon k}$ and $P_{\epsilon 0}$ in the Least squares sense.
6. Apply the transformation T: $P_{k+1} = T(P_0)$.
7. Compute the mean square distance $d_{k'}$ between $P_{\epsilon k+1}$ and $Y_{\epsilon k}$.
8. Terminate if the difference $d_k - d_{k'}$ is below a preset threshold or if the maximal number of iterations is exceeded, otherwise go to 2.

Figure 2(b) shows 4 out of 17 surface patches measured from different view points which were registered by the above described approach into the common coordinates system.

3.5 α-shapes to get a surface from a measured point cloud

A fair amount of work has been done to establish the definition of shape for a set of points in 2D as well as in 3D space [9, 10]. A mathematically rigorous definition was given by Edelsbrunner et al. [11]. The concept of *α-shapes* of a finite set of points for arbitrary real α was introduced in their article. The α-shape was derived from a straightforward generalization of a convex hull in two-dimensional space. Authors have also provided an optimal algorithm for the α-shape construction for sets of points in a plane. An α-shape is a concrete geometric object that is uniquely defined and can be exactly and efficiently computed. Lets imagine a plane colored by one background color. The points $p \in$

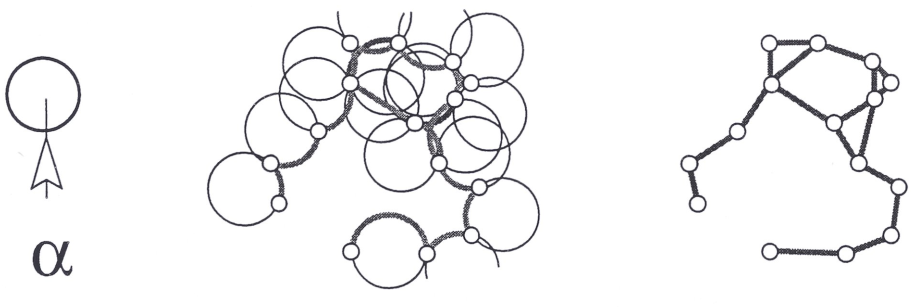

Fig. 3. (a) A ball of the radius α, (b) Points touched by an empty ball and connected by circular arcs, (c) In final α-shape, the arcs are replaced by the straight edges.

S are colored by a different color which distinguishes points from the background.

Now imagine eraser formed to the shape of circle with radius α, Fig. 3(b), which can be used to erase the background color, but non of the points p may be erased, Fig. 3(b). Before the erasing process starts the radius of the eraser can be adjusted to an arbitrary real $\alpha \geq 0$. During the actual procedure the radius α stays unchanged. The resulting object is called α-hull of the set of points S and it only depends on the points and the initial value of α, Fig. 3(c).

To make things more feasible for the computer representation, we straighten the surface of the object by substituting straight edges for the circular arcs. Final objects obtained after this substitution is the desired α-shape. It is a polyhedral complex which is not necessarily convex nor connected. For sufficiently large α, the α-shape is identical to the convex hull of S. As alpha decreases, the shape shrinks and gradually develops cavities. For sufficiently small α, the process of erasing results in an empty α-shape, which contains only the original points p. Thus the set of all real numbers α leads to a family of shapes capturing different levels of detail of the object defined by the particular point set.

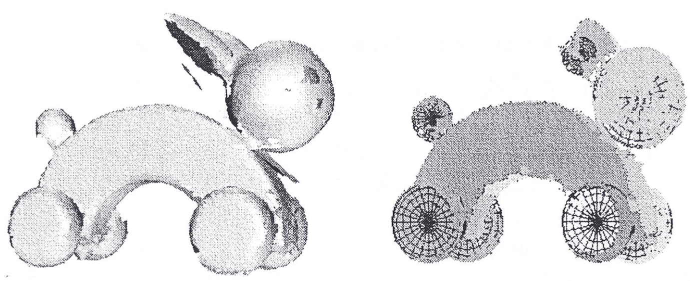

Fig. 4. (a)Registered α-shape, (b) Segmentated α-shape.

3.6 Segmentation

Artificial objects often consist of parts which can be modelled by a simple parametric models like planes, spheres, cylinders, etc. Ultimate goal of reverse engineering insists in recovering a CAD model of scanned objects, i.e in obtaining the parametric description of the parts. It is the task of *surface segmentation* to group the points which belong to the same part of the surface and to estimate its description.

We have exploited the approach to surface segmentation developed by Leonardis et al [12]. Their "Recover-and-Select" paradigm combines (i) the region growing with (ii) the description selection based on Minimum Description Length criterion to make the segmentation computationally feasible. In the phase of region growing, new points are being attached to the part if they can be described well by the model of the part. The selection, then, chooses an optimal subset of regions from all grown and often overlapping regions to end up with a partitioning of the points into the parts well described by the chosen models.

Figure 4(b) shows the result of α-shape segmented into the planar and spherical patches. Even the parts which cannot be well described by planes or spheres are modelled by some of the models. It therefore vital to include rich enough family of models into the set of expected shapes during the segmentation.

3.7 Beautified data sets

The result of α-shape is not a surface as it can contain singular points, edges, triangles and tetrahedra. If data are completely segmented, one can use parametric description to correct (beautify) the position of the points measured from the surfaces. It is possible to project the points onto their models or to the edges obtained by intersecting the adjacent models. Projected points are still connected by the edges defined in the phase of α-shape extraction. However, after the correction, the edges do not anymore form a tertrahedralization of 3D space. New edges have to be created by triangulating the corrected data inside the segmented surface parts. This process can, at least in the segmented parts, recover a surface from the α-shape.

Figure 5(a) shows α-shape coloured by the colour of the region into which the points were segmented. Thus, the wheels consist of a planar sides and spherical parts. The sides of the body have correctly been segmented into planes but cylindrical parts broke up into a few planar peaces as the cylinder model has not been in the set of expected models. The head and the tail are nicely modelled by spheres but the ears could not be well modelled by neither the plane not the sphere.

Figure 5(b) shows the corrected α-shape by projecting the vertices to their respective models and by removing the edges connecting different models. We can see, that the wheels are nicely recovered. Similarly, the head, the tail and the sides of the toy are correctly resolved. On the other hand, cylinders are incorrectly broken into the planes and ears are either left unmodelled or erroneously corrected to a sphere. Of course, correct results would be obtained if cylinders and ellipsoids were added to the model extraction during the segmentation.

4 Conclusions

We have demonstrated practical technology including new and improved methods for models construction from range images. The future work will aim at extending the beautifier to larger class of surface models, mainly to quadrics and superquadrics. We work on the methods, that use differential structures extracted from the surface to bootstrap the automatic range images registration [7].

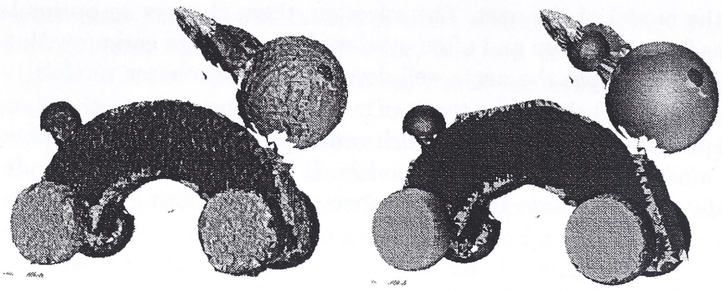

Fig. 5. (a) Segmentation, (b) Beautified mesh.

Acknowledgements

We benefit from discussions and help of doctoral students Pavel Krsek and Tomáš Svoboda. The segmentor used was created at the University of Ljubljana, Slovenia where mainly Aleš Leonardis and Bojan Kverh contributed. Beautifier has been programmed by Pavel Juran. Current version of the decimator has been implemented by Géza Kós from the Geometric Modelling Laboratory, MTA SZTAKI, Hungary.

This research was supported by the Czech Ministry of Education grant VS96049, the Grant Agency of the Czech Republic, grants 102/97/0480, 102/97/0855 and European Union grant Copernicus CP941068.

References

1. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 26(2):71–78, 1992.
2. K. Higuchi, M. Hebert, and K. Ikeuchi. Building 3-D models from unregistered range images. *Graphics Models and Image Processing*, 57(4):313–333, July 1995.
3. P. Uray. *From 3D point clouds to surface and volumes.* PhD thesis, Technische Universitaet Graz, Austria, October 1997.
4. P.J. Besl and N.D. McKay. A method for registration of 3-d shapes. *IEEE Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.
5. Tomáš Pajdla. Laser plane range finder - the implementation at the CVL. Technical Report Nr. K335-95-98, Czech Technical University, Prague, October 1995.
6. J. Rissanen. *Stochastic Complexity in Statistical Inquiry.* World Scientific Publishing Co., Series in Computer Science, IBM Almaden Research Center, San Jose, California U.S.A, 1989.
7. Pavel Krsek, Tomáš Pajdla, Václav Hlaváč, and Ralph Martin. Range image registration driven by a hierarchy of surface differencial features. In *Pattern Recognition and Medical Computer Vision 1998*, pages 175–183, Illmitz, Austria, May 1998. Osterreichische Computer Gesellschaft.
8. Tomáš Pajdla and Luc Van Gool. Matching of 3-D curves using semi-differential invariants. In *5th International Conference on Computer Vision*, pages 390–395, Cambridge, MA, June 1995. IEEE Computer Society Press.

9. R. A. Jarvis. Computing the shape hull of points in the plane. In *Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing*, pages 231–241. IEEE Computer Society Press, 1977.
10. J. Fairfield. Contoured shape generation: Forms that people see in dot patterns. In *Proceedings of the IEEE Conference on Cybernetics and Society*, pages 60–64. IEEE Press, 1979.
11. H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Transaction on Information Theory*, 29(4):551–559, 1983.
12. Bojan Kverh and Aleš Leonardis. Using Recovery-and-Select paradigm for simultaneous recovery of planes, second order surfaces and spheres from triangulated data. In *Proceedings of the Computer Vision Winter Workshop*, pages 26–37. IEEE Slovenia, February 1998.