Using Principal Curvatures and Darboux Frame to Recover 3D Geometric Primitives from Range Images

Eyal Hameiri, Ilan Shimshoni

Abstract— **We present a method for the recovery of partially occluded 3D geometric primitives from range images which might also include non-primitive objects. The method uses a technique for estimating the principal curvatures and Darboux frame from range images which we developed. After estimating the principal curvatures and the Darboux frames from the entire scene, a search for the known patterns of these features in geometric primitives, is performed. If a specific pattern is identified then the presence of the corresponding primitive is confirmed using these local features. The features are also used to recover the primitive's characteristics. The suggested application is very efficient since it combines the segmentation, classification and fitting processes, which are part of any recovery process, in a single process, which advances monotonously through the recovery procedure. The combined process does not use any least-squares fittings and therefore is very robust to inaccuracies in the local feature extraction process and to errors in segmentation. It was tested on series of real complex cluttered scenes, yielding accurate and robust recoveries of primitives.**

I. Introduction and Related Work

Three major processes are usually involved in tasks of primitive recovery from range images of a complex scene: segmentation, classification and fitting. During the segmentation an attempt to associate between a set of scene points and a common primitive or a patch of a primitive's surface is made. The classification process attempts to classify the primitive to a specific type of primitive. To complete the recovery, a fitting process attempts to fit each set of points within its type of primitives, as classified by the previous process, a unique primitive in a formulated representation. These three processes interact among themselves back and forth, iteratively. Recovery algorithms in 3D scenes differ by the methods they use for segmentation and fitting, and by the set of objects used for the classification. Our suggested recovery scheme is much more efficient since all three processes advance simultaneously and monotonously within a single process.

Two main approaches appear in the literature in regard to the segmentation of range images, the *edge-based* and *face-based* approach. The edge-based approach attempts to bound patches of surfaces, by revealing edge curves. It assumes that edge curves (moderate edges as well) are the borders between different types of surfaces which therefore have to be segmented separately. Such an approach is used in [9], [7]. On the other hand, the face-based approach attempts to group scene points which have special common features that indicate that the points belong to a common surface or object [1]. The segmentation component of our combined scheme is more related to the face-based approach but while all previous face-based methods start with seed regions that grow in an iterative procedure of hypotheses and validations, our approach segments the scene points simultaneously, not necessary as connective patches.

Fitting primitives to segmented areas has nearly always been based on least-squares fitting. Least-squares fitting is usually tailored to the specific type of surface being fit. Such fittings, can be found in [6], [8], [10]. Least-squares fitting procedure is performed in an iterative process which is strongly tied with the segmentation. Each proposed segmentation is validated using least-squares fitting. If the validation fails, the segmentation is modified accordingly and another validation is performed. Therefore, whatever type of fitting is made, bi-quadratic, B-spline or other least-squares fitting, these methods impose an inconvenient procedure. Least-squares fittings are also very sensitive to outliers and to noise in the range data. Our scheme for primitive recovery has an advantage in regard to these two difficulties described above, since it does not use any least-squares fittings at all.

Our approach for primitive recovery is based on estimating the principal curvatures and Darboux frame for each range point. Since the characteristics of the principal curvatures of geometric primitives are known for all types of primitives and distinguishable among the different types, primary segmentation and classification can be made by the principal curvatures values alone. But furthermore, from the values of the principal curvatures, associated with the segmented scene points, primary assumptions regarding the fitting is already being made as well. At this stage we use the local data stored in the Darboux frame of each scene point to refine the segmentation and classification and at the same time complete the fitting.

The main tool we use for recovering the type and parameters of the primitives is the Mean Shift Mode Estimator (M.S.M.E) [2], [3]. This tool enables us to recover the feature values, which are the modes of their distributions, without modelling this distribution (it does not need to be a Gaussian distribution) and it also deals very well with outliers in contrast with least-squares methods.

We tested our method on real scenes which contained

Eyal Hameiri is with the Dept. of Computer Science, The Technion - Israel Inst. of Technology, Haifa, Israel. E-mail: eyl1@cs.technion.ac.il.

Ilan Shimshoni is with the Dept. of Industrial Engineering and Management, The Technion - Israel Inst. of Technology, Haifa, Israel. E-mail: ilans@ie.technion.ac.il.

different types of primitives, in a variety of sizes and orientations, while some of them were partially occluded. The application proved to be very robust and accurate in its recoveries, even when clutter was added to the scene. Almost no false recovery occurred and as long as enough sampled scene points belong to a primitive, it was recovered quite accurately. Results of representative examples of these tests are presented towards the end of this paper.

II. Curvatures of Geometric Primitives

In this section we shall examine in short (not analytically but more visually) the characteristics of principal curvatures and directions for several geometric primitives, and ways to image the principal curvatures from an entire object or scene.

A. Principal Curvatures Image

Let us assume that we gather the principal curvatures from a finite number of surface points. We can describe all these values within a two dimensional and discrete histogram in which one axis represents the minimal principal curvature (in this work we use the vertical axis) and the other axis represents the maximal principal curvature. The histogram value at a specific location reflects the amount of minimal vs. maximal principal curvature values found in the data. The histogram can also be displayed as a gray-level image. Thus, for every object or surface an image of its principal curvatures can be generated. This image is invariant to rigid transformations, as long as the same portion of the surface is analyzed, but it can be altered due to partial occlusion of the object.

B. Principal Curvature Images of Primitives

We will now review in short the principal curvature images of several primitive objects.

Planes

In the trivial case of planes, all directional curvatures equal 0, thus for each point on a plane $\kappa_1 = \kappa_2 = 0$. The principal directions in planes are therefore undefined. All planes have a common gray level image of their principal curvatures - one bright point at the center of the image (Figure 1).

Spheres

For a sphere, the absolute value of each principle curvature is the inverse of its radius. The sign of the curvatures is negative due to our definition of the orientation of the normals. The principal directions are undefined. The curvatures image is one bright point on the negative side of the main diagonal at a distance of the inverse of the radius of the sphere from both negative axes (see Figure 1).

Cylinders

The cylinder's bases are considered as planes. As for the rest of the cylinder's points, they all share the same values for the minimal and maximal principal curvatures. The maximal principal curvature, $\kappa_1 = 0$, and its related tangent direction, T_1, aligns with the main axis of the cylinder. The minimal principal curvature, $\kappa_2 = -1/r$, where r is the cylinder's radius. The direction related to the minimal principal curvature, T_2, is the tangent to the circular normal section at P. The principal curvatures image is one bright point on the negative vertical axis at a distance of $1/r$ from the origin (Figure 1).

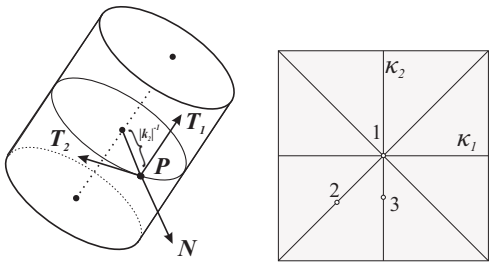

(a) (b)

Fig. 1. (a) Cylinder - principal curvatures and Darboux frame; (b) Principal curvature typical image: 1 - plane, 2 - sphere, 3 - cylinder.

Cones

All points lying on a cone (except for its apex and base) share the same maximal principal curvature - $\kappa_1 = 0$. T_1, is always pointing from P to the apex (see Figure 2). T_2, is the tangent to the elliptical normal section contained in the plane which is tilted by the opening angle of the cone, α, relative to the plane containing the circular cross section that passes through P. If r denotes the cone radius at P then the minimal principal curvature is given by $-r^{-1} \cdot \cos\alpha$.

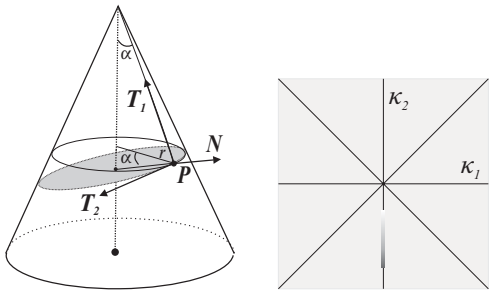

(a) (b)

Fig. 2. Cone: (a) Darboux frame; (b) Typical image of principal curvatures.

The image of the principal curvatures is theoretically a half infinite and straight line along the negative vertical axis which starts at $(0, -R^{-1} \cdot \cos\alpha)$ where R is the radius of the base. Practically, if we look at a finite number of data points then the image is a finite straight line from $(0, -R^{-1} \cdot \cos\alpha)$ to $(0, -\widehat{R}^{-1} \cdot \cos\alpha)$ (see Figure 2)where $\widehat{R}$ is the cone's radius at the closest analyzed point to the cone's apex.

Tori

At all points on a torus, the minimal principal curvature equals $-r^{-1}$ where r is the minor radius of the

torus. The minimal direction at point P is the torus tangent at P that is also contained in the plane which is perpendicular to the center circle of the torus (see Figure 3). If R denotes the major radius then the maximal principal curvature at P varies from $(R-r)^{-1}$ to $-(R+r)^{-1}$. The corresponding principal direction is the radial tangent. Analytically, the principal curvatures are:

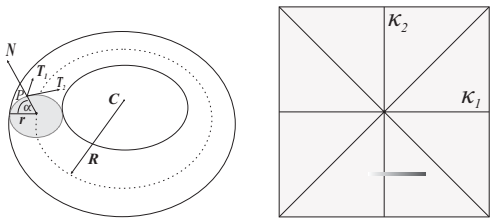

Fig. 3. Torus: (a) Darboux frame; (b) Typical image of principal curvatures.

$$\kappa_P^1 = -\frac{\cos(\alpha)}{R+r\cos(\alpha)} \qquad \kappa_P^2 = -\frac{1}{r}$$

The principal curvature image is a horizontal straight line from $(-(R+r)^{-1}, -r^{-1})$ to $((R-r)^{-1}, -r^{-1})$ crossing the negative vertical axis at $(0, -r^{-1})$ (see Figure 3).

III. The recovery strategy

As was described in Section II, every type of primitive has its unique "signature" in the principal curvatures histogram, of the entire scene. For most types of primitives there is even a different signature for primitives of different dimensions within the same type. Any part of a geometric primitive has the same pattern of signature as the entire primitive has. Therefore, it is possible to discover the presence and characteristics of a specific geometric primitive by analyzing the curvatures histogram of the entire scene, even if the object is partially occluded. This stage of the process, in which we examine the principal curvatures histogram, will be termed as the first stage of the recovery process. At the end of the first stage, hypotheses for the presence of specific types of primitives, or even for specific primitives, already emerge.

The second stage of the recovery process involves the Darboux frames which were extracted from the scene. All the Darboux frames of a specific primitive inherit the same change in their orientations, which is of course determined by the orientation of the primitive. Therefore, whenever an initial hypothesis for the presence of a specific primitive emerges from the curvatures histogram, the Darboux frames at the hypothetical primitive points can be used to confirm or reject the hypothesis. The second stage of the recovery is different for each type of primitive and will be discussed in detail later.

Curvature histograms do not preserve the geometric shapes of the scene's objects, i.e. geometrical proximity is not necessarily translated to proximity in the curvatures histogram. The histograms, also do not have the ability to separate the points of non-primitive objects from the points of the primitives because general free-form objects might contribute to the histogram at the location which indicates a primitive. Another problem is that due to noise in the data, a very noisy sample of the scene might yield a contribution to the histogram at a location which substantially differs from its correct location. In other words, noise in the data is translated in to noise in the histograms. Moreover, because curvatures are differential features of the second degree, the errors in the curvatures histogram are quite substantial.

The difficulty is to find, in a noisy curvatures histogram, or in what will be defined later, as "multidimensional histograms", that also contains noise, the signature of a specific primitive. This signature might be blurred due to noise in the input data or due to non-primitive objects. We address this problem by using the *Mean Shift Mode Estimator*, described in Section IV-A, to search for peaks in histograms. Curvature signatures of geometric primitives, as appear in the gray level image of the curvatures, take the form of dots and lines at very specific locations. Since general free-form surfaces do not concentrate their impact at certain points or along certain lines, signatures of primitives will still be detectable over a local bias level of the curvatures histogram. Therefore, a search for local peaks and series of local peaks will still reveal the primitives, even with the presence of non-primitive objects. Looking for local peaks also reduces the effect of noise in the curvatures histogram. Errors in the estimated curvature values are translated in to errors in the positions in the histogram. Therefore by locating local peaks in histograms we are neutralizing the effect of noise as well as ignoring outliers. In the second stage of the recovery process we also look for peaks in histograms of other features than curvatures like principal directions, normals and combinations of them. The advantages of the search for peaks in histograms, in regard to the ability to isolate the signature of a certain primitive, are also relevant when the histogram is based on other features, as will be explained individually for each type of primitive, in Section IV-B.

IV. The implementation

We begin with the estimations of normals, principal curvatures and principal directions for each point in the scene as explained in [4], [5]. The recovery procedure is performed by a quick search for peaks in n dimensional (nD) histograms of the features and/or of their functions. The first stage involves only the 2D histogram of the principal curvatures and the second stage uses histograms of higher dimensions as will be detailed later.

A. Mean Shift Mode Estimator

In order to search for peaks in an nD histogram we use the nD version of a Mean Shift Mode Estimator - M.S.M.E. [2], [3]. We define an nD window centered around a kernel point in the nD histogram. Then, iteratively until convergence is achieved, the center of mass within the points covered by the window is calculated and the window is centered over the new center of mass. M.S.M.E in one dimension is illustrated in Figure 4.

Two main factors control and optimize the M.S.M.E

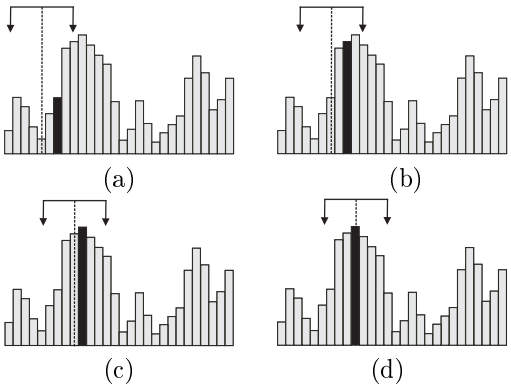

Fig. 4. One Dimensional Mean Shift Mode Estimator. The dark bars indicate the new center of mass: (a) (b) and (c) New center of mass is computed within the window's coverage and the window is centered over the new location; (d) The new center of mass is located at the previous one's location. The procedure has converged.

functionality - the set of initial points where the window is first located and the window's "width" (might be multidimensional). Too many initial points will add extra unnecessary computations and not enough of them will leave unrevealed local peaks. The window's "width" controls the M.S.M.E resolution, i.e. the ability to distinguish one local peak from another "adjacent" local peak. A wide window might cause two or more adjacent local peaks to be detected as one at a location where there isn't one at all. A narrow window might reveal unwanted extra peaks, which are caused by noise, and located between two or more local peaks of higher values. A narrow window might also detect a single local peak, which due to noise in data splits into two or more adjacent local peaks, as several separate local peaks. In general, the M.S.M.E. tends to converge very rapidly and has the advantage of robustly detecting peaks of unknown distributions contaminated by outliers.

B. The recovery by the type of primitive

The second stage of the recovery deals with each type of primitive individually. We search for peaks over multidimensional histograms, different histograms for each type of primitive. The main idea of this stage is to detect features like curvatures, vectors of the Darboux frames or a combination of them which are common to a group of scene points which are assumed, as a result of the first stage, to lie on a common type of primitive. Therefore, these multidimensional histograms involve only 3D points suspected to lie on that specific type of primitive. A scheme of the entire recovery procedure is presented in Figure 5.

B.1 Planes.

A peak in the 2D curvatures histograms at a point were $\kappa_1 = \kappa_2 = 0$ indicates one or more planar elements. Let $N \cdot P + D = 0$ be the equation of the plane where N is the plane's normal, P is a point on the plane and D a scalar. All points on a specific plane share the same surface normal. Another common feature is $D = -N \cdot P$. Now, if a planar element is indeed present in the analyzed scene there must be a set of scene points which share the same values for N and for $-N \cdot P$. We therefore expect to find a local peak in the histogram of N and $-N \cdot P$. The histogram is generated from scene points which were suspected by the first stage to lie on planes. This way we are reducing the amount of computations and also blocking the possibility that arbitrary scene points that do not lie on planes might create, a local peak in the histogram of N and $-N \cdot P$. When we locate a local peak we actually get a confirmation for the presence of a planar element and also reveal its parameters according to the specific location of the local peak. Thus by locating one or more local peaks we can separate between different planes and find for each one of them its parameters.

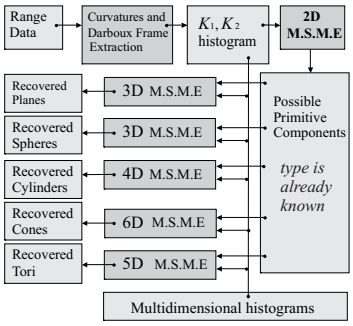

Fig. 5. Scheme of the recovery procedure

B.2 Spheres.

Spherical elements appear in the curvatures histogram as peaks along the main diagonal where $\kappa_1 = \kappa_2 \neq 0$. Each peak represents one or more spherical element of a distinguishable radius and therefore it is more efficient to analyze each peak separately. A peak at (κ, κ) suggests that one or more spherical element of radius $|\kappa^{-1}|$ exists in the scene. We can recover their characteristics and distinguish between them by recovering their centers. For every

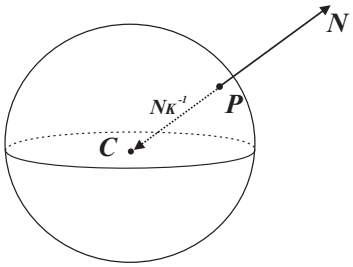

Fig. 6. The common features in the sphere case

point P lying on a sphere, the sphere's center is located at a distance of the radius, i.e. $|\kappa^{-1}|$, in the direction opposite to the normal N at P (see Figure 6). Thus we search for peaks in the 3D histogram of $P - |\kappa^{-1}|N$. In addition to getting a conformation to the presence of spherical elements of radius $|\kappa^{-1}|$, we also recover their centers. Note that if the first stage of the recovery, locates a false local peak which was generated from arbitrary scene points that do not lie on a common sphere (or primitive in the general case), then the second stage will not confirm the hypoth-

esis because these points will fail to generate a local peak in the histogram of the second stage.

B.3 Cylinders.

A local peak in the curvatures histogram located at $\kappa_1 = 0$, $\kappa_2 \neq 0$ is an indication for the existence of one or more cylinders of radius $|\kappa_2{}^{-1}|$. We deal with each "cylinder peak" separately, but still each peak might indicate several cylinders of the same radius but of different orientations and/or locations within the scene.

We use the direction of the cylinder's main axis and a point on one of the main planes through which the main axis passes as common features. Let P be a 3D scene point contributing to the peak in the histogram. The main axis direction is given by the maximal direction at P - T_1 (see Figure 7).

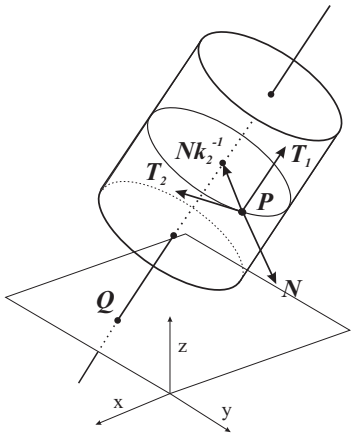

Fig. 7. The common features in the cylinder case

A point on the cylinder's main axis can be found at a distance of the radius - $|\kappa_2{}^{-1}|$ from P in the direction opposite to the normal, N. Now we can find the intersection points of the axis with the main planes (at least one exists). We find the first intersection point with the $X \circ Y$, $Y \circ Z$ and $Y \circ Z$ planes, in that order. Keeping that order guarantees that for all points which lie on a common cylinder we will get the same intersection point. Therefore a local peak in the 4D histogram of the two features will be found confirming the existence of the hypothesized cylinders as well as distinguishing between them and revealing the equations of their main axis.

B.4 Cones.

A series of local peaks along the negative vertical axis of the curvatures histogram, is an indication for the presence of one or more cones. It is not possible at this stage to distinguish one cone from another. Furthermore, since cylinder peaks are located on the negative vertical axis as well, their associated scene points might be included within the input for our cone revealing steps, where only there, these points will be identified as not associated with a cone. Thus, all scene points associated with local peaks along the negative vertical axis of the curvatures histogram, and only them, are considered as "suspected cone points".

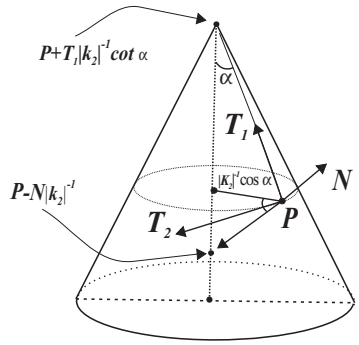

Fig. 8. The common features in the cone case.

The common characteristics we use here are the cone apex and its main axis direction. As illustrated in Figure 8 the maximal principal direction - T_1, always points at the apex. $|\kappa_2|^{-1}\cos(\alpha)$ is the value of the cone's radius at point P, lying on it, where α is the apex angle. Thus, the apex is located at $P+T_1|\kappa_2|^{-1}\cot(\alpha)$. In addition, another point can be found on the cone's main axis in a distance of $|\kappa_2|^{-1}$ from P, at the opposite direction of the normal to the cone's surface, at P. The direction of the main axis is obtained by subtracting this second point from the apex yielding $T_1\cot(\alpha) + N$ as the common axis direction.

We end with a 6D histogram in which one of the dimensions is the angle α. Since the angle has not been recovered yet, each scene point suspected to be a "cone point" will contribute to the histogram along a series of values of α.

B.5 Tori.

Series of local peaks, in the curvatures histogram, which lie on horizontal lines that cross the negative vertical axis, are primary indications for tori in the scene. We can already distinguish between different possible sets of tori, of different minor radii for each set. The minor radius is $|\kappa_2|^{-1}$, where κ_2 is the minimal principal curvature for all points related with such series of peaks.

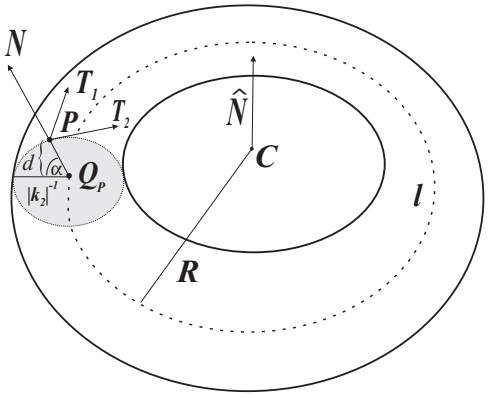

Fig. 9. The common features in the torus case

We split the second stage into three sub-stages. The scene points which are used as input for each sub-stage are those who are related to a single local peak found in the previous sub-stage.

Let l denote the circle which passes through the torus' center (see Figure 9) and $\hat{N}$ the normal to the main plane of the torus (not a unit length normal). A point Q_P on l can be found at distance $|\kappa_2|^{-1}$ from a point P on the torus, and in the opposite direction of N, i.e. $Q_P = P - |\kappa_2|^{-1}N$. The maximal principal direction is parallel to the main plane and therefore $T_1 \cdot \hat{N} = 0$.

Let the equation of the torus' main plane be $\hat{N} \cdot Q + 1 = 0$ where Q is a point on the plane. Q_P lies on the main plane and therefore $\hat{N} \cdot Q_P + 1 = 0$.

Let us denote $\hat{N}$, T_1 and Q_P by their components:

$$\hat{N} = \left(\hat{N}_x, \hat{N}_y, \hat{N}_z\right), \qquad T_1 = \left(T_x^1, T_y^1, T_z^1\right)$$
$$Q_P = \left(Q_x^P, Q_y^P, Q_z^P\right)$$

By using the above facts we can obtain that

$$\left(T_x^1 Q_y^P - T_y^1 Q_x^P\right) \hat{N}_y + \left(T_x^1 Q_z^P - T_z^1 Q_x^P\right) \hat{N}_z + T_x^1 = 0$$

Thus, for each scene point suspected to lie on a torus with a minor radius of $|\kappa_2|^{-1}$ we can "draw" a line over the two dimensional histogram of $\hat{N}_y$ and $\hat{N}_z$. Looking for local peaks over this histogram will reveal couples of coordinates for $\hat{N}_y$ and $\hat{N}_z$ whenever a torus with the above parameters is present in the scene. Each such peak defines a possible main plane of one or more tori.

Now, we have to distinguish between tori of different major radii and/or locations, while their minor radius and main planes are already known. The maximal principal curvature at a point P on a torus is given by $\kappa_1 = -\frac{\cos(\alpha)}{R + r\cos(\alpha)}$. The distance from P to the main plane denoted by d (see Figure 9) can be found. Therefore, the angle α is available and the major radius can be extracted. Thus, a 1D histogram for the major radius is generated followed by a search for local peaks that reveal all possible major radii.

The input for the last sub-stage is a set of scene points, suspected to lie on a common torus with known, hypothetical values for its minor and major radii and for its main plane. To complete the recovery we need to find the torus' center, denoted by C. As we saw for each point P a related point Q_P which lies on the central circle can be found at $P - |\kappa_2|^{-1}N$. As illustrated in Figure 10 for each point P

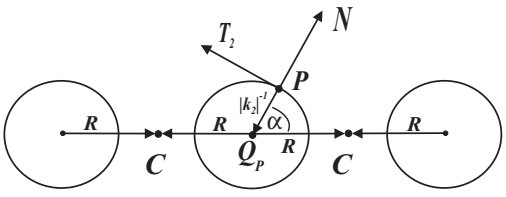

Fig. 10. Possible locations of tori centers.

there are two possible locations for the torus' center. Actually, this ambiguity is an indication for two different possible tori. But, a local peak over the histogram of the centers will be generated only if a torus is indeed located at that center. To locate the centers which correspond to P we first move to Q_P. The centers are located at a distance of the hypothetical major radius, denoted by R, from Q_P. Since both of them lie on the intersection line between the torus main plane and the plane defined by N and T_2, they can be determined. Now, as we locate local peaks over the centers histogram, we can confirm or contradict the presence of each specific hypothetical torus, in the scene. In case of confirmation, the location of the local peak reveals the torus center and therefore the torus is fully recovered.

C. Several Details of the Implementation

Two main problems emerge when using the M.S.M.E. in the first and second stages - the window's width and how to choose the kernel points. An optimal window's resolution along one of its dimensions should be somehow related to the range of values along that dimension. When the dimension involves principal curvatures and since most of the extreme values of curvatures are due to noise or to restricted areas of the scene (like sharp edges) we do not take the entire range of values to be the basis for the window's width. Instead, we empirically found that optimal values for the window's dimensions can be fixed as a percentage (usually, 2 to 5 percents) of the range within which 90% of the values of the relevant dimension exist. Therefore, in each dimension, according to the relevant histogram, a different width for the window is used. The parameters for this criterion may vary when the level of noise changes or when the scene contains many sharp edges. If the dimension involve only 3D points, vectors or angles we take a fraction of the entire range.

The second issue which has to be addressed is how to pick the initial points for the M.S.M.E process? We choose this set of kernel points by building a discrete version of the histogram. The resolution at each discrete dimension is of the same size as was assigned to the width of the window in this dimension. Otherwise, if the resolution is higher (smaller intervals) a repeated search over the same locations will be performed, while lower resolution might cause local peaks to be missed. Every point on the discrete supporting plane of the histogram, over which the height crosses a fixed threshold, is taken as a kernel point. We fixed these thresholds as a small percentage (0.1% to 0.5%) of the scene points which generate the relevant histogram. Note that discrete histograms are used only for choosing an optimal set of kernel points and not during the search for local peaks in the histogram, where we use the continuous data set. Taking all these steps and measures guarantees quick and efficient detection of the peaks in a histogram although there might be cases of false alarms due to extreme noise, or cases in which local peaks are missed because of inadequate scanned area of a primitive, due to occlusion.

V. Experimental Results

We have tested our application on a large number of range images acquired by the Cyberware Laser Scanner, Model 3030 (1993 model), in different scenes of partially occluded geometric primitives and non-primitive objects.

Our application succeeded to recover accurately every primitive as long as a minimal number of its points (relatively to other objects or other primitives of its type) were sampled in the range image.

We present here two typical results of running our algorithm. The first scene, consists of two identical spherical objects placed on the same horizontal plane, two different cylinders in different orientations and one box which was placed such that three of its facets are visible (see Figure 11). The ground truth of the cylinders and spheres radii was obtained by physical measurements. The error of these measurements is $\pm 0.01cm$. Quantitative recovered results

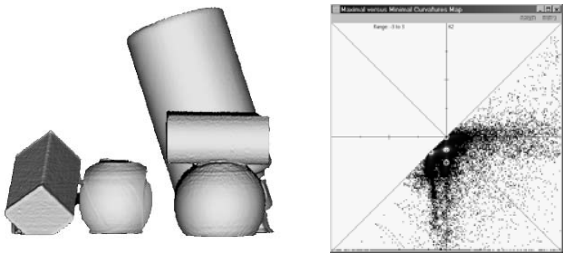

(a) (b)

Fig. 11. Real scene of partially occluded primitives: (a) Range image; (b) Local peaks at the end of the first stage illustrated upon the center of the principal curvatures image.

are summarized in Table I: for radii of primitives, when it is relevant, a comparison of the recovered to the measured values is presented. The units are in centimeters. "O" stands for a Cylinder's orientation, "P" for a point on a cylinder's main axis, "N" and "D" for a plane's normal and the free coefficient in its equation respectively, and "C" for a center of a sphere. At the end of the first stage, four local peaks were located at (-0.394332, -0.394332) with 6669 scene points, (-0.003446, -0.334160) with 14,843 points, (0.000462, -0.677247) with 4068 points and (-0.002089, -0.003453) with 2484 scene points (see Figure 11). The four local peaks reflect the scene as expected. All primitives were detected and there were no false alarms. The box was recovered as 3 planes. Although we did not make any accurate measurements of the features of the scene, all the checks we have performed agreed with the recovery results. For example, the inner product of the normals of the three recovered planes are 0.012438, 0.000241 and 0.000965 - as expected since the planes are orthogonal to each other. The centers of the two identical spherical elements have almost the same value for their vertical coordinate. Again, this was expected since they were placed at the same vertical height. The smaller cylinder's main axis is practically horizontal and indeed this was its orientation as can be seen in the scene's illustration.

The second presented scene contains free-form objects together with several primitives (see Figure 12). We placed in this scene two different cylinders (with different radii and orientations), one box, one sphere and two general and non-primitive objects. The configuration of the objects creates partial occlusion of some of the primitives. In figure 13 the M.S.M.E process is demostrated for the spherical object in the scene.

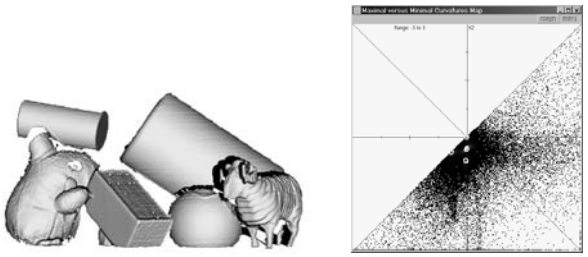

(a) (b)

Fig. 12. Real scene of partially occluded primitives and free-form objects: (a) Range image; (b) Local peaks at the end of the first stage illustrated upon the center of the principal curvatures image.

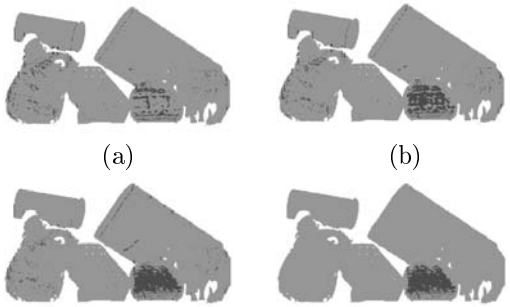

(a) (b) (c) (d)

Fig. 13. Segmentation along the first and second stages which is related to a single local peak of the curvatures histogram of the second scene. (a) (b) and (c) illustrate the segmentation along the first stage; (d) Segmentation at the second stage. Note that in the second stage all points not on the sphere dissappear.

The results are summarized in Table II. All primitives were detected, and there were no false recoveries. The inner product between the normals of the 3 orthogonal plane's are 0.00214, 0.00039 and 0.01062. The visible base of one of the cylinders was recovered as an additional plane and indeed its normal is very close to the direction of cylinder's axis.

VI. Discussion

We have presented a robust and accurate method for the recovery of 3D geometric primitives from complex cluttered range scenes. The application presents a new recovery approach which does not use any least-squares fitting. It maintains accurate recoveries even when the range images include partially occluded objects or contain free-form objects as well. The application was tested on range images of various scenes. All primitives, in all scenes, were recognized and their parameters were recovered very accurately. Almost no false recoveries were obtained and the only few cases in which a primitive was not detected were when the object was not sampled at enough surface points, due to extreme occlusion. The recovery application also demonstrates the possibilities in using our curvature and Darboux frame estimators within a practical application.

TABLE I

SCENE WITH PRIMITIVES ONLY - RECOVERED RESULTS VS. MEASURED FEATURES.

Cylinder a	radius	O_x	O_y	O_z	P_x	P_y	P_z
Measured	2.98	–	–	–	–	–	–
Recovered	2.993	-0.303	0.952	-0.028	3.459	0.000	-3.995
Cylinder b	radius	O_x	O_y	O_z	P_x	P_y	P_z
Measured	1.41	–	–	–	–	–	–
Recovered	1.477	0.999	-0.010	-0.014	0.000	-0.960	1.277
Sphere a	radius	C_x	C_y	C_z			
Measured	2.52	–	–	–			
Recovered	2.536	-1.572	-4.556	-2.768			
Sphere b	radius	C_x	C_y	C_z			
Measured	2.52	–	–	–			
Recovered	2.536	4.195	-4.529	1.387			
Plane a		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		-0.637	0.661	0.396	-2.107		
Plane b		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		-0.018	-0.540	0.841	-4.180		
Plane c		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		0.770	0.529	0.355	5.067		

TABLE II

A SCENE OF PRIMITIVES AND FREE-FORM OBJECTS - RECOVERED RESULTS VS. MEASURED FEATURES.

Cylinder a	radius	O_x	O_y	O_z	P_x	P_y	P_z
Measured	2.98	–	–	–	–	–	–
Recovered	3.070	0.757	-0.650	-0.009	0.000	2.662	-4.205
Cylinder b	radius	O_x	O_y	O_z	P_x	P_y	P_z
Measured	1.41	–	–	–	–	–	–
Recovered	1.443	0.899	-0.320	0.297	0.000	0.992	-1.878
Sphere	radius	C_x	C_y	C_z			
Measured	2.52	–	–	–			
Recovered	2.575	3.374	-3.251	0.443			
Plane a		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		0.343	0.825	0.447	1.451		
Plane b		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		-0.803	0.014	0.596	-2.795		
Plane c		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		0.494	-0.564	0.661	-3.594		
Plane d		N_x	N_y	N_z	D		
Measured		–	–	–	–		
Recovered		0.914	-0.291	0.282	4.215		

REFERENCES

[1] Paul J. Besl and Ramesh C. Jain. Segmentation Through Variable-Order Surface Fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(2):167–192, 1988.

[2] Y. Cheng. Mean shift, mode seeking and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, August 1995.

[3] D. Comaniciu and P. Meer. Distribution free decomposition of multivariate data. *Pattern Analysis and Applications.*, 2(1):22–30, 1999.

[4] E. Hameiri. Estimating the principal curvatures and the darboux frame from real 3D range data and its application to the recovery of primitives. Master's thesis, The Technion, 2001.

[5] E. Hameiri and I. Shimshoni. Estimating the principal curvatures and the darboux frame from real 3D range data. In *Symp. on 3D Data Processing Visualization Transmission*, 2002.

[6] D. Marshall, G. Lukacs, and R. Martin. Robust segmentation of primitives from range data in the presence of geometric degeneracy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(3):304–314, March 2001.

[7] M. J. Milroy, C. Bradley, and G. W. Vickers. Segmentation of a wrap-around model using an active contour. *Computer-aided Design*, 29(4):299–320, 1997.

[8] D. Rogers and N. Fog. Constrained B-spline curve and surface fitting. *Computer Aided Design*, 21(10):641–648, 1989.

[9] David R. Smith and Takeo Kanade. Autonomous scene description with range imagery. *Computer Vision, Graphics, and Image Processing*, 31(3):322–334, September 1985.

[10] Ma Weiyin and J. P. Kruth. Parameterization of randomly measured points for least squares fitting of B-spline curves and surfaces. *Computer-aided Design*, 27(9):663–675, 1995.