Geometrical Cloning of 3D Objects via Simultaneous Registration of Multiple Range Images

Peter J. Neugebauer
Fraunhofer Institute for Computer Graphics
Wilhelminenstr. 7, 64283 Darmstadt, Germany
neugeb@igd.fhg.de, http://www.igd.fhg.de/www/igd-a7

Abstract

In this paper, we present a method for the reconstruction of real world objects from multiple range images. One major contribution of our approach is the simultaneous registration of all range images acquired from different scanner views. Thus, registration errors are not accumulated, and it is even possible to reconstruct large objects from an arbitrary number of small range images. The registration process is based on a least-squares approach where a distance metric between the overlapping range images is minimized. A resolution hierarchy accelerates the registration substantially. After registration, a volumetric model of the object is carved out. This step is based on the idea that no part of the object can lie between the measured surface and the camera of the scanner. With the marching cube algorithm a polygonal representation is generated. The accuracy of this polygonal mesh is improved by moving the vertices of the mesh onto the surface implicitly defined by the registered range images.

1. Introduction

3D scanners providing range images are able to scan the shape of an object from one view. However, it is not possible to scan the complete object at once due to topological and geometrical limitations. First, parts of the 3D-object are occluded or may lie in shadow. Second, the object might be significantly larger than the scanner is able to capture in one view. In order to simplify the scanning procedure, we must assume that the relative orientations between the different scanner views are unknown. Thus, several unregistered range images showing only partial views of the object must be registered in order to calculate the unknown relative orientations.

The reconstruction of objects is divided in the steps registration of all range images and integration of all range images in one complete geometrical model. The registration is a highly nonlinear problem and hence an initial estimation of the relative orientation is required. Interactively selected point correspondences are used in our approach to compute the initial estimate. For the reconstruction of large objects from several small range images, it is very important that the registration errors are not accumulated. The unique feature of our approach is that we solve this problem by a simultaneous registration of all range images. Normally, the simultaneous registration of numerous range images requires a great deal of computation time. With our resolution hierarchy, the computation time of the iterative registration is significantly reduced. With our statistical termination criterion, the registration is iterated until all information pertaining to the actual resolution is exploited, and then we switch to the next finer resolution. Thus only few iterations are necessary in the highest resolution.

For the generation of a model, it is very important to handle self-occlusions and scanning errors. In section 7, we develop a visibility criterion which is the basis for detecting self-occlusions and scan errors. The test is simply based on the idea that a point in 3d space lying between the camera and the surface cannot belong to the object. By applying this test to all voxels of a volume, it is easy to sculpture a volumetric model of the object. The volume describes the topology of the object and approximates its shape within the desired level of detail. By finding isosurfaces in the volumetric model a polygonal representation is generated. As an alternative to the sculpturing approach, the object can be rendered directly. The direct rendering is performed by establishing a modified z-buffering algorithm which makes use of the visibility criterion. With our approach, it is possible to reconstruct concave and convex objects, and even objects with holes out of an arbitrary number of range images.

Reconstruction has various applications in computer graphics and virtual reality. The geometry of real world objects can be documented, e.g. statues before and after restoration, or designer models can be reconstructed in the rapid prototyping process. With methods well-known in

computer graphics, the appearance of the reconstructed objects can be changed or simulated in different environments.

2. Previous works & Drawbacks

The registration of range images is based mainly upon two approaches: relational matching and minimization techniques. Relational matching [19] [6] techniques try to find correspondences of previously segmented image primitives such as significant points, edges or planar surfaces. Under the enormous number of possible assignments, the correct matching is searched with tree search methods. Then, the relative orientations of the range images are computed from the corresponding image primitives. In order to make the registration robust and to limit the computation time, the number of image primitives should be small. Thus, the calculated relative orientation is stressed with a certain inaccuracy.

In order to obtain a better registration accuracy, minimization techniques are basically better for yielding the relative orientation between the range images more accurately. Thereby, a distance metric between two range images is defined and minimized as a function of the unknown location. In general, this is a nonlinear problem. An initial estimate of the unknown location must be computed with other methods. In [2], the distance metric is based on point correspondences. The *iterated-closest-point* algorithm finds the nearest positions on one surface to a collection of points on the other surface and then transforms one surface so as to minimize the distance metric. In [14] and [3], the distance metric is based on correspondences between points on the surface and nearby tangent planes on the other surface. A rigid motion is found that minimizes the point-to-tangent collection directly and then iterates.

The problem of building complete models from a number of range images is considered by Turk and Levoy [18]. The integration and registration of the range images is done step-by-step. In all stages of the algorithm, the range images are represented as triangulated meshes. A modified closest-point algorithm is used to register the range images. Jump edges and self-occlusions are detected when the distance between neighbored vertices of the mesh exceeds a certain threshold. This heuristic criterion requires that the uncertainties of the distance values are smaller than the spacing between the samples. To avoid accumulating the registration errors, a cylindrical range image of the complete object must be captured, and then all other ranges are aligned to this cylindrical view. The main drawback of this approach is that the scanner must be able to capture a cylindrical view of the complete object.

In [1], Bergevin, Laurendeau and Poussart present an approach where the partial range views are represented as triangulated meshes. Features characterizing the triangles are used for a relational matching technique. The meshes are represented in different resolutions to reduce the number of potential correspondences. Multiple trials about correct matching are followed and the best are selected for the generation of better assumptions. The initially found locations are improved by a minimization technique. It is uncertain how many assumptions are necessary to generate the correct result. The results demonstrate that, even for small objects, many initial assumptions must be followed.

In [7], a feature-based approach based on the curvature of the surfaces is also used. The images are mapped on a spherical image, and the transformations between the views are computed by matching the spherical images. In this approach, the objects may not contain holes. For large objects, it is not possible to map a significant part of all range images onto the sphere.

Assuming that all acquired range images are already registered Curless and Levoy present in [5] an volumetric approach to reconstruct complete objects. The range images are converted to a distance function stored in a volumetric array. The distance functions are combined for all range images. By extracting isosurfaces from the volume grid the object is reconstructed. They apply a hole filling algorithm to interpolate unseen portions of the surface which will otherwise appear as holes in the reconstruction.

None of the known approaches does a simultaneous registration of all range images. Thus, there is a great need for the registration algorithm presented in section 5.2. In contrast to other approaches, our approach does not need heuristic criterions for the generation of unambiguously reconstructed objects.

3. The processing pipeline

In our approach, the reconstruction process is divided into several steps (see figure 1). First, the object is scanned from several directions with a 3D scanner. Then, the acquired range images are segmented with threshold based methods. We discriminate 3 different classes: object, background and undefined regions (shadows, artifacts). It should be noted that nowadays most of the 3D scanners are able to detect shadows automatically, and, that a segmentation into object and background is normally not needed. Therefore, the segmentation only requires little interaction, if any. However, in some limited cases, it is helpful to manually mask regions in the range images which do not contain helpful information for the reconstruction process.

After the segmentation, the relative orientations of all range images are computed in two steps. In the first step, interactively marked corresponding points define a rough estimate of the required orientation. In the second step, the exact relative orientation is determined automatically with a least-squares surface registration with sub-pixel accuracy.

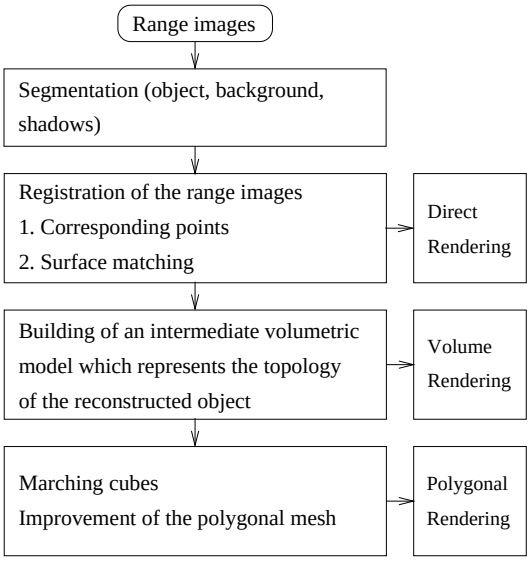

Fig. 1. The processing pipeline.

All range images are mapped into a common world coordinate system in order to sculpture a 3D volume which approximates the shape the object. The marching cube algorithm is applied to resample the object geometry in 3D space. The output of the marching cube is improved by automatically moving the vertices of the mesh onto the surface defined by the original range images. As an alternative, the reconstructed object can be rendered directly by a modified z-buffering algorithm.

4. 3D scanning

In the first step, the object is acquired from several directions using a range image scanner. For our statue example the range images are captured for example with a structured light range finder which can be easily adjusted for different viewing directions. Each relevant part of the object surface should be visible in at least one view. In order to allow for a registration, neighboring views should show some overlay; typically 20% — 50% overlap is sufficient. Shadows are detected automatically by the 3D scanner. A segmentation of the object from the background can be easily done by threshold-based methods.

5. Registration of multiple range images

5.1. Point-based registration

The range images are successively integrated into the model step-by-step. Starting with the first image, the user selects a new image to be integrated next. The user has to manually identify at least 3 points in the new image and the 3 corresponding points in any of the previously selected images. The accuracy of this point identification procedure need only be very crude: the subsequent steps still work with orientational errors of 30 degrees, or even more. The rough relative orientation between the new image and the world coordinate system is defined with at least three corresponding points between the new image and all previously integrated images. The detection of the relative orientation with this approach is described e.g. in [6].

The rough estimate of the relative orientation is improved by applying the algorithm in section 5.2. This process is repeated with each new range image. Assuming that the relative orientations of all previously integrated range images are already known, the registration algorithm described in section 5.2 can be easily modified to run much faster. After integrating the last range image, the described algorithm in section 5.2 should be performed as described to get the best results for the relative orientations of all range images.

5.2. Simultaneous registration of range images

The registration of multiple range images $S_0, \ldots, S_{n-1}$ is a minimization problem where a distance measurement between the overlapping surface portions is minimized. In our approach, the *registration of all range images is done simultaneously in order not to accumulate registration errors*. For this reason, the relative orientations of all range images must be roughly known initially.

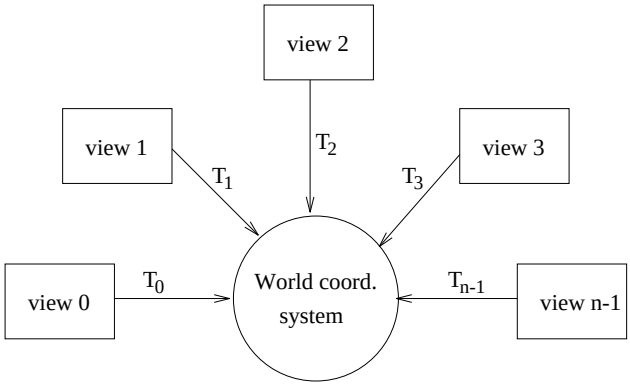

Fig. 2. The relative orientation for each range image with respect to a world coordinate system.

The relative orientation $\vec{\theta}_i$ for each image S_i is defined according to a common world coordinate system. The orientation vector $\vec{\theta}_i$ consists of 6 parameters: 3 rotation angles (the Euler angles) and 3 translation components. The rotation matrix $\mathcal{R}_i^{(\theta)}$ and the translation vector $\vec{t}_i^{(\theta)}$ are derived directly from the parameter vector $\vec{\theta}_i$. The rotation and

translation of points $\vec{x} \in S_i$ to the world coordinate system are defined with the operator $T_i^{(\theta)}$.

$$T_i^{(\theta)}\vec{x} \stackrel{\text{def}}{=} \mathcal{R}_i^{(\theta)}\vec{x} + \vec{t}_i^{(\theta)} \tag{1}$$

Now, we define a matching distance $D(S_i, S_j)$ by measuring the distance between the rotated and translated surfaces S_i and S_j. Let us call the surface S_i the model and the surface S_j the scene. This matching distance has to be minimized as a function of the relative orientations $\vec{\theta}^{\,\mathrm{T}} = \left(\vec{\theta}_0^{\,\mathrm{T}}, \ldots, \vec{\theta}_{n-1}^{\,\mathrm{T}}\right)$. After minimization, the rotated and translated surfaces should become almost identical in the overlapping areas, i.e. the total residual error ε^2 should become minimal.

$$\varepsilon^2 = \min_{\vec{\theta}} \sum_{i \neq j} D(T_i^{(\theta)} S_i, T_j^{(\theta)} S_j) \tag{2}$$

Since the rotation matrix depends nonlinear on the rotation angles and the range images itself are considered as arbitrary functions over the scanner grid, the formulated optimization problem is a highly nonlinear problem and cannot be solved analytically. Thus, it must be solved numerically with correction methods.

An estimation of the relative orientation $\vec{\tau}$
is iteratively improved by a correction vector $\vec{\delta}$
until the minimum of the least-squares problem is found

Fig. 3. The principle of iterative correction.

As shown in figure 3, correction methods start with an initial estimate which is iteratively improved by a correction vector $\vec{\delta}$. In our problem, the unknown mappings $T_i^{(\theta)}$ must be split into a known part $T_i^{(\tau)}$ and a correction $T_i^{(\delta)}$ according to $T_i^{(\theta)} = T_i^{(\delta)} T_i^{(\tau)}$.

$$\varepsilon^2 = \min_{\vec{\delta}} \sum_{i \neq j} \varepsilon_{ij}^2 \tag{3}$$

$$\varepsilon_{ij}^2 = D(T_i^{(\delta)} T_i^{(\tau)} S_i, T_j^{(\delta)} T_j^{(\tau)} S_j) \tag{4}$$

The correction vector $\vec{\delta}$ is searched in the local neighborhood of $\vec{\tau}$ by linearizing the problem at the estimated solution $\vec{\tau}$. For this reason, the model S_i appearing in the matching distance is represented piecewise by its tangential planes $\vec{n}_{ik}(\vec{x} - \vec{x}_{ik}) = 0$ where k is the index of the available model points. The normal vectors $\vec{n}_{ik}$ are simply computed with local gradient operators in the range image S_i at the measured point $\vec{x}_{ik}$. $T_i^{(\tau)}$ can be effectively applied by rotating and translating the tangential plane:

$$\vec{n}_{ik}^{\,(\tau)} = \mathcal{R}_i^{(\tau)} \vec{n}_{ik} \qquad \vec{x}_{ik}^{\,(\tau)} = T_i^{(\tau)} \vec{x}_{ik} \tag{5}$$

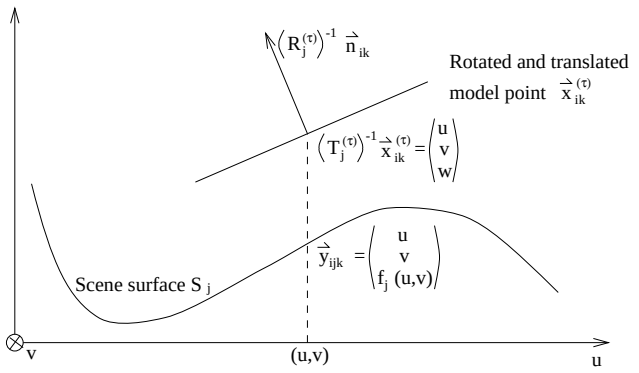

Fig. 4. Projecting the model point to the scene yields the corresponding scene point.

A corresponding scene point $\vec{y}_{ijk}^{\,(\tau)}$ can be easily found by mapping the point $\vec{x}_{ik}^{\,(\tau)}$ into the coordinate system of the range image S_j and then projecting it onto the parameter grid of the image S_j (see Fig. 4). In the parameter grid of the scene image S_j, the depth is interpolated and the interpolated point $\vec{y}_{ijk}$ has to be mapped back to the world coordinate system.

$$\vec{y}_{ijk}^{\,(\tau)} = T_j^{(\tau)} \vec{y}_{ijk} \tag{6}$$

For most of the points, the distance between the corresponding points $\vec{x}_{ik}^{\,(\tau)}$ and $\vec{y}_{ijk}^{\,(\tau)}$ should be small. Exceptions to this rule can be detected with the outlier test in section 6.4. The outliers have to be rejected from the registration algorithm. With these definitions, the distance between the model and the scene is expressed as

$$\varepsilon_{ij}^2 = \sum_k \varepsilon_{ijk}^2 \tag{7}$$

$$\varepsilon_{ijk} = \mathcal{R}_i^{(\delta)} \vec{n}_{ik}^{\,(\tau)} \left(\left(\mathcal{R}_j^{(\delta)} \vec{y}_{ijk}^{\,(\tau)} + \vec{t}_j^{\,(\delta)} \right) - \left(\mathcal{R}_i^{(\delta)} \vec{x}_{ik}^{\,(\tau)} + \vec{t}_i^{\,(\delta)} \right) \right) \tag{8}$$

After linearization and some algebraic manipulations, we obtain for each point k a $1 \times 6n$ matrix $\mathcal{A}_{ijk}$ and a scalar s_{ijk}. With these variables the total residual ε^2 can be expressed as

$$\varepsilon^2 = \min_{\vec{\delta}} \sum_{i \neq j,k} \varepsilon_{ijk}^2 = \min_{\vec{\delta}} \sum_{i \neq j,k} \| \mathcal{A}_{ijk} \vec{\delta} - s_{ijk} \|^2 \tag{9}$$

with

$$s_{ijk} = \vec{n}_{ik}^{\,(\tau)} \cdot (\vec{x}_{ik}^{\,(\tau)} - \vec{y}_{ijk}^{\,(\tau)}) \tag{10}$$

$$A_{ijk} = \left(\underbrace{0 \ldots 0}_{6i \times 1} \underbrace{C_{ijk}}_{6 \times 1} \underbrace{0 \ldots 0}_{6(l-i-1) \times 1} \right) +$$

$$\left(\underbrace{0 \ \dots \ 0}_{6j \times 1} \ \underbrace{-C_{ijk}}_{6 \times 1} \ \underbrace{0 \ \dots \ 0}_{6(l-j-1) \times 1} \right) \tag{11}$$

$$C_{ijk} = \begin{pmatrix} \vec{n}_{ik}^{(\tau)} \times \vec{y}_{ijk}^{(\tau)} \\ -\vec{n}_{ik}^{(\tau)} \end{pmatrix}^{\mathrm{T}} \tag{12}$$

6. Solving the nonlinear least-squares problem

Since the world coordinate system can be freely defined, there is a linear dependency between the unknown parameters. This dependency can be avoided by defining that the relative orientation of the first range image is identical to the orientation of the world coordinate system. Based on this assumption, the derived linearized equation system can be easily reduced. The remaining problem can be solved with standard methods. With the Newton-Taylor algorithm, the correction vector $\vec{\delta}$ is computed with

$$\vec{\delta} = \left(\sum_{i \neq j,k} \mathcal{A}_{ijk}^{\mathrm{T}} \mathcal{A}_{ijk} \right)^{-1} \sum_{i \neq j,k} (\mathcal{A}_{ijk}^{\mathrm{T}} s_{ijk}) \tag{13}$$

In fact, we use the Levenberg-Marquardt method [12] [10] to solve the numerical problem. It differs from the Newton-Taylor method in two points. First, the expected magnitude of all unknowns is made equal by substituting the unknowns with scaled ones. The square of good scaling factors are found on the main diagonal of the normal matrix $\mathcal{N} = \sum_{i \neq j,k} \mathcal{A}_{ijk}^{\mathrm{T}} \mathcal{A}_{ijk}$. Second, a multiple of the unit matrix is added to the normal matrix. With a suitable chosen multiple of the unit matrix, it can be guaranteed that ε^2 decreases in each iteration and that numerical problems never occur.

6.1. Statistical considerations

The statistical regression model allows an estimation of the mean deviation σ_i of the components $\vec{\delta}$. A reasonable assumption is that the residuals ε_{ijk} are uncorrelated normal random variables with mean 0 and variance σ^2. A reliable estimate of σ^2 is given by the normalized total residual error. In order to get a reliable estimate, the denominator of the quotient has to be decreased by the number of degrees of freedom. If m specifies the number of all equation rows, σ^2 is given by:

$$\sigma^2 = \varepsilon^2 / (m - 6n) \tag{14}$$

The corresponding covariance $\Sigma_{\delta\delta}$ can be computed effectively in (15):

$$\Sigma_{\delta\delta} = \sigma^2 \left(\sum_{i \neq j,k} \mathcal{A}_{ijk}^{\mathrm{T}} \mathcal{A}_{ijk} \right)^{-1} \tag{15}$$

Then, the corresponding variances σ_i^2 are found on the main diagonal of the covariance matrix.

6.2. Termination criterion

In each iterative algorithm, the termination criterion plays an important role. The algorithm may terminate if all components of the correction vector $\vec{\delta}$ are lower than a predefined threshold. We developed a better method in our approach. Our test is based on a statistical test for the hypothesis $\|\vec{\delta}\| = 0$. The algorithm terminates if the correction has no statistical relevance and is only caused by noise.

Therefore, we need a significance test which verifies the hypothesis $\|\vec{\delta}\| = 0$. Under the assumption that the residuals ε_{ijk} (see 9) are normal random variables and that the hypothesis $\|\vec{\delta}\| = 0$ is true, the variable t in (16) has a χ^2_{6n}-distribution. For the test, the components of the correction vector $\vec{\delta}$ are normalized by their expected inaccuracies:

$$t = \sum_i \frac{\delta_i^2}{\sigma_i^2} \tag{16}$$

With $t < 6n$, we have a statistical test for the above hypothesis. If the test is true, no more iterations of the nonlinear least-squares problem are needed. In practice, we need at least $6m + 20$ points to gain statistically reliable results for the test variable t.

6.3. Resolution hierarchy

The robustness and the efficiency of the proposed algorithm can be improved by introducing a resolution hierarchy. In the first stages of the interactive process, only a few points of the model (according to the notation in section 5.2) are used. In subsequent stages, additional model points are included in the computation process. In practice, the chosen points in the lower resolutions should be distributed uniformly in the image plane. Thus the resolution hierarchy can be simply applied just by choosing only those points in in the image grid whose grid coordinates are a power of 2^d. However, the corresponding scene points should be found in the full resolution of the range images. Each level in the resolution hierarchy is maintained until the statistical termination criterion is fulfilled. Then, the resolution is doubled in each dimension. In this way, the algorithm quickly gains the full resolution of the 3D model. In the finest resolution, the computed location parameters reach their maximal accuracy. The speed-up of this simple resolution hierarchy is a factor of 20 – 100 because the run time of the registration

depends mainly on how often the registration algorithm is iterated in the highest resolution.

6.4. Elimination of outliers

Missing parts, occlusions and scan errors influence the accuracy of the registration algorithm. Especially in the first iterations, the initial estimate of the relative orientation is very inaccurate. Thus, the rotated and translated surfaces do not fit very well together. By eliminating outliers, most of the occluded or missing parts are automatically detected. With the same procedure, another important problem which arises during the registration can be solved. All pairs of model points $\vec{x}_{ik}^{\,(\tau)}$ and their corresponding scene points $\vec{y}_{ijk}^{\,(\tau)}$ which do not belong to each other must be rejected from the registration.

Outliers of the model points cause a residual ε_{ijk} which lies above the average of the other residuals. Normally, the average of the residuals are computed at the end of each iteration. Fortunately, there is a very efficient way to detect outliers earlier. In each iteration of the applied correction method, the total residual and the mean residual σ of the last iteration is known. s_{ijk} (see 10) specifies the residual belonging to the previous iteration. Thus, outliers are detected with the test

$$|s_{ijk}| > 3\sigma_{\text{last}} \tag{17}$$

However, in practice, it is preferable to perform additional outlier tests, e.g. a test depending on the directions of the normal vectors or, simply, the following test:

$$\left|\vec{x}_{ik}^{\,(\tau)} - \vec{y}_{ijk}^{\,(\tau)}\right| > 3\sigma_{\text{last}} \tag{18}$$

These tests are very efficient, because the values of the variables must be computed in the registration algorithm anyway. If some model points are eliminated in one iteration, the same points can be taken into account in the next iterations. Thus, this test is robust with respect to faulty decision-making in previous iterations. Statistical considerations show that the outlier test does not influence the convergence of the registration algorithm. This can be shown by considering the iterated outlier test. The number of eliminated outliers converges very fast. In the case of normal distributed residuals, approximately 1 percent of the model points are eliminated in the infinitely often iterated outlier test.

6.5. Improvements

The registration of multiple range images is a highly nonlinear problem with several small local minima besides the global minimum. Hence, the iterative algorithm of the nonlinear least-squares problem can converge in a small local minimum. The risk of this behavior can be reduced by smoothing the total residual as a function of the unknown parameters. As a result, small local minima disappear. This smoothing can be performed by averaging the normal vectors of the model in their local neighborhood. Of course, such a smoothing influences the accuracy of the found location parameters. Hence, the smoothing should be applied in the first iterations. In subsequent stages, the smoothing should be reduced successively.

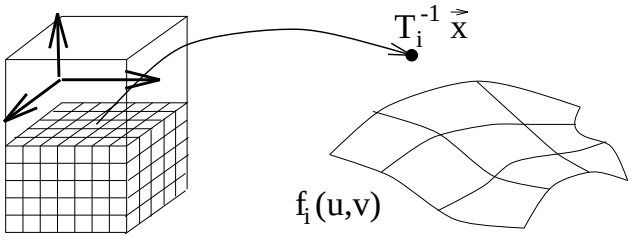

Fig. 5. Carving out the volume.

7. Visibility criterion

In contrast to other approaches where the range images are typically considered as a set of points, we consider the range images as partially defined surfaces. All points below this surface are invisible from the scanner, whereas points lying between the surface and the scanner are visible, and, thus, definitely do not belong to the object. This idea is used to formulate a visibility criterion (see 20 and figure 5) which is very useful for the detection of occluded surfaces. With the projection operators P_x, P_y and P_z

$$P_x\ \vec{x} \stackrel{\text{def}}{=} \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}^{\mathrm{T}}\ \vec{x} \tag{19}$$

the projection to the base vectors of the coordinate grid are defined. For simplicity, let us assume here that the scanner is far away from the object. Then, a point $\vec{x} = \begin{pmatrix} u & v & w \end{pmatrix}^{\mathrm{T}}$ given in world coordinates does not belong to the object, if

$$\exists i:\ \ P_z\ T_i^{-1}\ \vec{x}\ >\ f_i(P_x\ T_i^{-1}\ \vec{x}, P_y\ T_i^{-1}\ \vec{x}) \tag{20}$$

$f_i(u,v)$ stands for the interpolated depth at the grid coordinates (u,v) in the range image S_i. Now, the visibility criterion is used to reconstruct the object unambiguously even in the case of complicated topologies.

After registration and estimation of the relative orientation, the partial views of the range images are used to carve out an intermediate volumetric model of the object. This is done by applying the visibility criterion in equation 20 for each voxel of the volume. A voxel does not belong to the

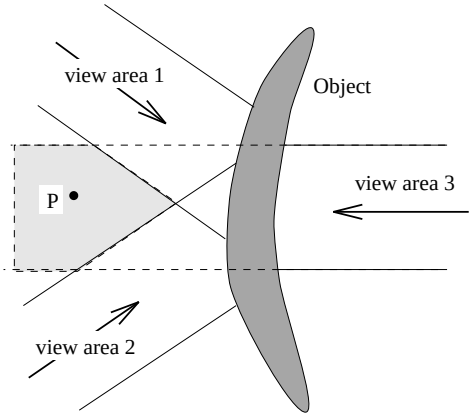

Fig. 6. Undefined points in 3D space.

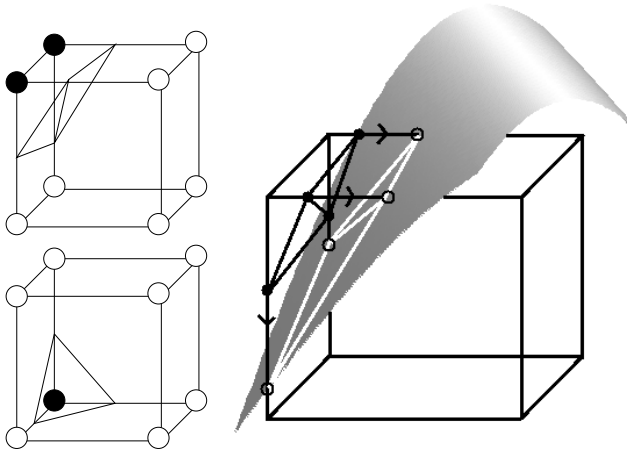

Fig. 7. Triangulation of a cube by applying Marching-Cubes and improving the mesh.

object if it is visible or if none of the range images are defined in criterion 20.

In some cases, for some voxels lying far away from the object, it is not possible to directly decide whether they belong to the object. Figure 6 illustrates this. If the overlapping scan views (view 1 and 2) are captured under a certain angle, points far away from the object may lie outside the scan views. Assume that there is another scan view (scan view 3) capturing the back face of the object. Then, the point P is not visible in the sense of criterion 20. The criterion 20 alone cannot decide if P belongs to the object. This problem is solved with a post-processing step on the volume. All connected voxel volumes have to be detected in the volume. Then, all volume regions without connections to a voxel definitely belonging to the object can be deleted. Voxels belonging to the object can be detected by mapping the measured points in the range images back into the volume grid. In addition to the deletion step, closed, hollow spaces in the volume are completely filled. The intermediate volume describes the topology of the object and approximates its shape within the desired level of detail. It accelerates the reconstruction process and allows the reconstruction of objects without any topological constraints.

With this intermediate volume, the visibility criterion can be completed. A point P in 3D space is called visible if the criterion 20 is true. If the criterion 20 is false, or if all range images in the criterion are undefined for the requested point, then the intermediate volume must decide the visibility. If the point P falls into a cube of the volume where the majority of the bordering voxels vote for visibility, then the point P is said to be visible; otherwise, it is marked as invisible. Normally, all of the bordering voxels have an equal status. If not, this tells us that some parts of the surface are not captured with the range finder, or the overlap of the scanning views is smaller than the size of one cube in the intermediate volume.

8. Reconstruction and rendering

8.1. Volume rendering

The intermediate volume can be rendered directly by using a volume rendering technique [16] [15]. Nowadays, these techniques are quite fast. An example of a rendered volume can be seen in figure 8. Drawbacks are the large amount of voxels necessary to represent the volume in high quality. Typical volumes showing a sufficient level of detail consist of 100 — 700 slices. In combination with other developments (e.g. [4]) this step could be improved.

8.2. Polygonal rendering

After the generation of the volumetric model, the marching cube algorithm [11] with a look-up table that resolves ambiguous cases [13] (see figure 7) can be applied to generate a polygonal representation. The accuracy of this polygonal mesh is improved by moving the vertices of the mesh onto the surface implicitly defined by the registered range images. The point is moved between the voxels of the cube until the visibility criterion in 20 degenerates to an equation. The exact coordinates of the moved vertices are interpolated in the range images. This step is similar to a ray casting algorithm which is combined with the visibility criterion. This way, the full accuracy of the scanning device is exploited at all vertices of the mesh, and concave object shapes are also modeled correctly. The amount of polygons can be significantly reduced by applying a polygon reduction [17] [9]. However, the computational complexity of this polygonal

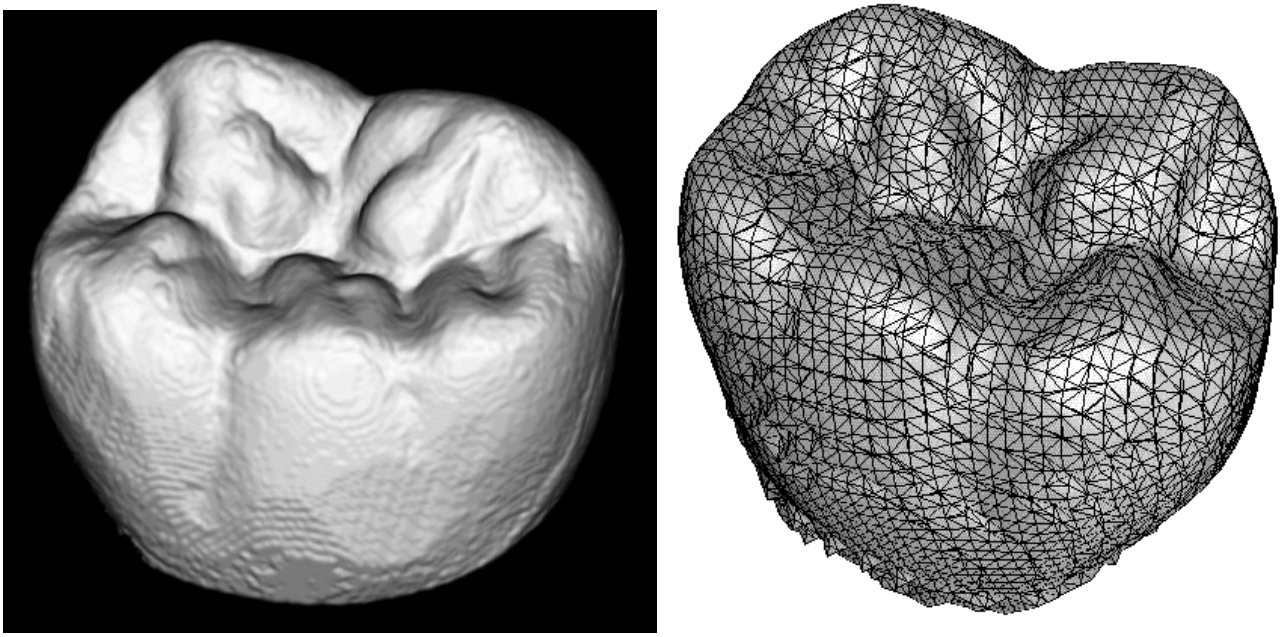

Fig. 8. Steps through the reconstruction process: volumetric model, polygonal model.

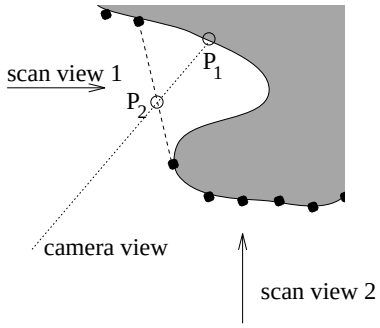

Fig. 9. Using the visibility criterion for detecting self-occlusions.

reduction is very high. An alternative is the direct rendering of the object which is described in the next section.

8.3. Direct rendering

According to the view of a virtual camera, the object can be rendered with a modified z-buffer algorithm. The z-buffer algorithm must be modified in order to avoid problems arising from self-occlusions and scanning errors. In figure 9, such a situation is illustrated. From the view of the camera P_1 is interpolated in range image 1 and P_2 in range image 2. Due to self-occlusions, P_2 is not placed on the surface of the object. This can be easily detected with the visibility criterion in section 7. Therefore, only those points which are not rejected by the visibility criterion should be used for rendering. The criterion must be evaluated before the rendered point is compared with the z-buffer. It should be noted that for large objects, the direct rendering yields the best rendering results, because the full accuracy of the scanning device can be exploited.

9. Results

Results are shown for a molar tooth. The range images of the tooth are captured with a laser scanner based on the principle of triangulation. The scanning of one view (256.000 points) takes about 45 seconds. The distance between two scanned points is 25 μm and the height resolution is 10 μm. In order to obtain the necessary accuracy, the triangulation angle is quite large. Consequently, there are obviously a lot of shadow areas. In figure 8 the volumetric model and a low resolution polygonal model are shown. The obtained CAD data is rendered with methods well-known in computer graphics and the resulting image is shown in figure 12. The tooth was reconstructed out of 25 range images acquired from different scanning views. The interactive reconstruction steps are performed within a few minutes, say 20 seconds per range image. The simultaneously registration takes approximately half an hour on a SPARC-station 10. The run times for sculpturing and post-processing the volume are

Object	scans	volume size	time/min
Statue	180	$200 \times 74 \times 77$	10
Statue	180	$500 \times 185 \times 194$	70
Statue	180	$700 \times 299 \times 272$	210
Tooth	25	$170 \times 200 \times 193$	15

Table 1. Sculpturing and post-processing the volume

shown in table 1.

Another example is presented of a statue $182cm$ tall. The statue was scanned with a structured light scanner which can be easily adjusted for different viewing angles. Each range image shows a window of approximately $40 \times 40cm^2$. The statue has a complicated topology, because it contains holes and many occlusions. Due to its size and complicated geometry, the statue was scanned from 180 views from all directions (see figure 11). All range views are combined with the described method. This takes approximately one day. The overall registration error is $0.5mm$. The original and the reconstructed statue are shown in figure 10. Some range images and the reconstruction program are shown in figure 11.

10. Conclusions and future work

One main contribution of our approach lies in the simultaneous and robust registration of all range images with subpixel accuracy. This way, it is even possible to reconstruct large objects from numerous small range images. The registration process is based on a least-squares approach where a distance metric between the overlapping views is minimized. A resolution hierarchy accelerates the registration process. The registration process is robust even in the presence of small artifacts that typically appear in range images.

With the developed visibility criterion, it is possible to decide whether or not each point in 3D space belongs to the object. This knowledge is exploited for the generation of a model or simply by rendering the reconstructed object directly. The object can be resampled by generating an intermediate volumetric model. The generated volumetric model can be rendered directly or it can be used for the generation of a polygonal mesh. The resolution of the volume can be chosen in accordance to the desired accuracy of the tessellation.

Further work has to be done to generate an adaptive volume resolution and to generate a polygonal mesh which approximates the original surface in nearly the same quality with less polygons. Splitting the volume into subvolumes may solve this problem.

References

[1] R. Bergevin, D. Laurendeau, and D. Poussart. Registering range views of multipart objects. *Computer Vision and Image Understanding*, 61(1):1–16, January 1995.

[2] P. J. Besl and N. D. McKay. A Method of Registration of 3-D Shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.

[3] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.

[4] C. Connolly. Cumulative Generation of Octree Models from Range Data. In *Proc. 1st IEEE International Conf. on Robotics and Automation*, page 25, 1984.

[5] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH '96 Conference*, pages 303–312, New Orleans, Louisiana, 1996.

[6] O. D. Faugeras and M. Hebert. The Representation, Recognition, and Locating of 3-D Objects. *The International Journal of Robotics Research*, 5(3):27–53, 1986.

[7] K. Higuchi, M. Hebert, and K. Ikeuchi. Building 3-d models from unregistered range images. *Graphical Models and Image Processing*, 57(4):315–333, July 1995.

[8] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *SIGGRAPH '92 Conference*, pages 71–78, Chicago, 1992.

[9] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. In *SIGGRAPH '93 Conference*, pages 19–26, Anaheim, California, 1993.

[10] C. L. Lawson, Richard J. Hanson. *Solving Least Squares Problems*. Prentice-Hall, Englewood Cliffs N.J., 1974.

[11] W. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *Computer & Graphics*, 21(4), 1987.

[12] D. W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *J. Soc. Indust. Appl. Math.*, 11(2):431–441, 1963.

[13] C. Montani, R. Scateni, and R. Scopigno. A modified lookup table for implicit disambiguation of marching cubes. *Visual Computer*, 10(6), 1994.

[14] P. J. Neugebauer. Hochgenaue Objektlokalisation in Tiefenbildern. In *Visualisierung — Rolle von Interaktivität und Echtzeit*, Sankt Augustin, Schloß Birlinghoven, June 1992. GMD — Gesellschaft für Mathematik und Datenverarbeitung mbH.

[15] G. Sakas. Fast rendering of arbitrarily distributed volume densities. In *Proceedings of Eurographics'90*, pages 519–530. Elsevier Science Publishers B.V., North–Holland, September 1990.

[16] G. Sakas and S. Walter. Extracting surfaces from fuzzy 3d-ultrasound data. In *SIGGRAPH '95 Conference*, pages 465–474, Los Angeles, California, 1995.

[17] W. Schroeder, J. Zarge, and W. Lorensen. Decimation of triangle meshes. *Computer Graphics (SIGGRAGPH '92 Proceedings)*, 26(2):65–70, July 1992.

[18] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *SIGGRAPH '94 Conference*, pages 311–318, Orlando, Florida, 1994.

[19] G. Vosselman. *Relational Matching*. Springer-Verlag, 1992.

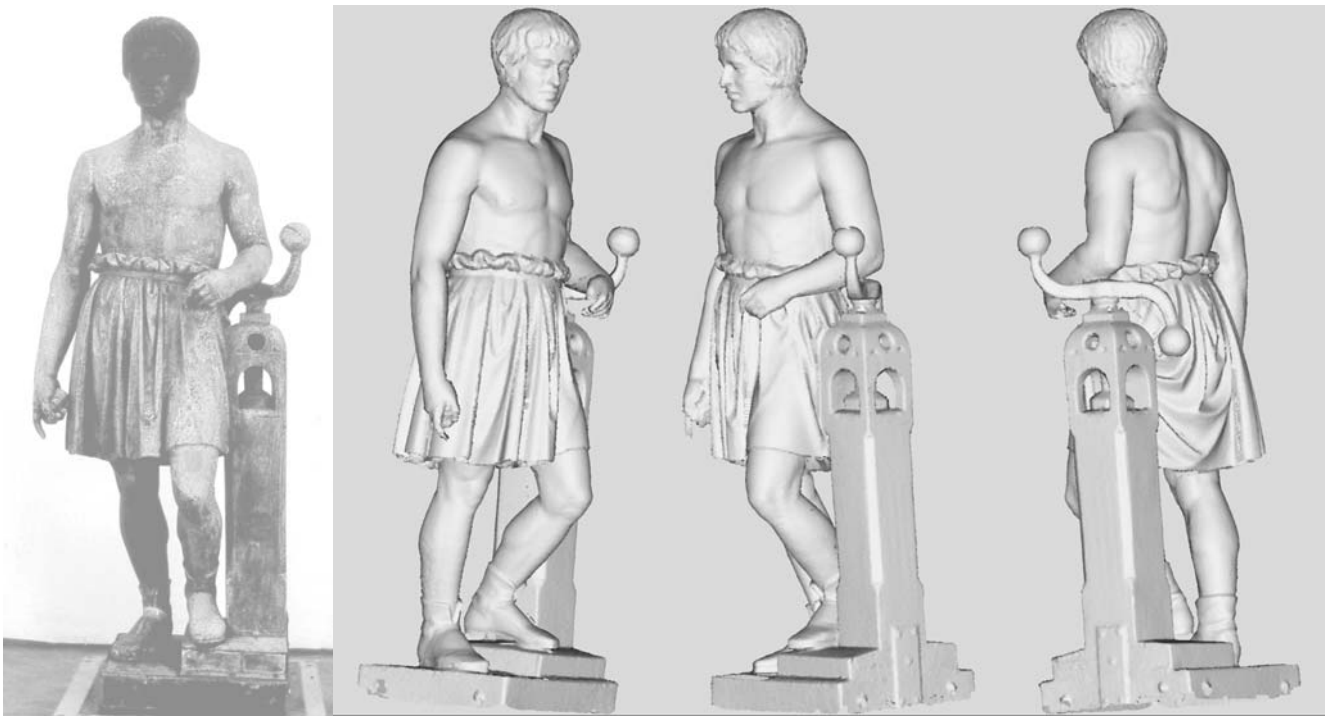

Fig. 10. Left: The original statue. Right: Direct rendering of the reconstructed statue.

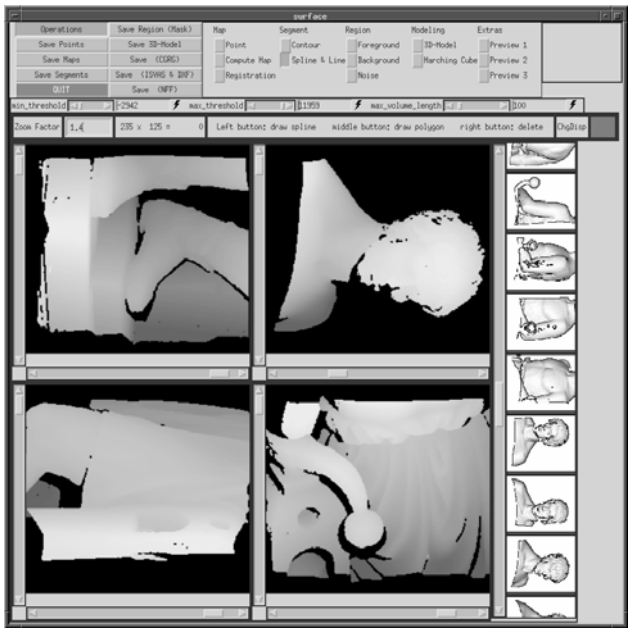

Fig. 11. Modeling a statue from 180 views.

Fig. 12. A reconstructed molar tooth.