This paper appears in: The Second International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT 2004)

Specularity Elimination in Range Sensing for Accurate 3D Modeling of Specular Objects

Johnny Park and Avinash C. Kak
Robot Vision Lab, Purdue University, West Lafayette, IN. USA
{jpark,kak}@purdue.edu

Abstract

We present a novel range sensing method that is capable of constructing accurate 3D models of specular objects. Our method utilizes a new range imaging concept called multi-peak range imaging, which accounts for the effects of mutual reflections. False measurements generated by mutual reflections are then eliminated by applying a series of constraint tests based on local smoothness, global coordinate consistency and visibility consistency. We show the usefulness of our method by applying the method to three real objects with specular surfaces. The ground truth data for those three objects were also acquired in order to evaluate the elimination of false measurements and to justify the selection of the parameters in the constraint tests. Experimental results indicate that our method significantly improves upon the traditional methods for constructing reliable 3D models of specular objects with complex shapes.

1. Introduction

In recent years, there has been considerable interest in constructing accurate three-dimensional models of real-world objects for applications where the focus is primarily on visualization of the objects by humans. This interest is fueled by the recent technological advances in range sensors, and the rapid increase of computing power that now enables a computer to represent an object surface by millions of polygons and that allows such representations to be visualized interactively in real-time. Researchers have shown that the state-of-the-art techniques can now construct detailed 3D models of objects ranging from small figurines to large statues. Although they have established the feasibility of constructing accurate 3D models, there still remain several challenging issues. One of these challenging issues arises from the fact that many objects have surface materials that are not ideal for range sensors. Various surface properties that cause difficulties in range imaging include specular surfaces, highly absorptive surfaces, translucent surfaces, and transparent surfaces. Some researchers have tried to simply do away with such surface-related problems by painting the object or coating the objects with removable powder to ensure that the surfaces reflect the light source diffusely. Obviously, this approach is not desirable and may not even be feasible for real-world objects outside the laboratory.

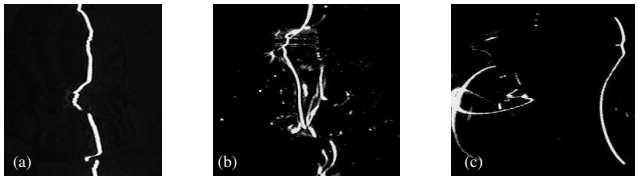

Figure 1.
Camera images of 3 different objects when a laser stripe is projected onto the objects. (a): Object with a diffuse surface. Notice that only a single peak for each camera scan line exists. (b),(c): Objects with specular surfaces. Due to the effects of mutual reflection, multiple peaks in the same scan lines exist.

Of the various surface-related properties we mentioned above, surface specularity is one of the more problematic material properties. Specularity causes mutual reflections that give rise to ghosts in the measured structured-light data. Depending on the extent of specularity, the presence of these ghosts can make it difficult to localize a data point that corresponds to the object point that was actually illuminated. Figure 1 shows camera images of three different objects when a laser stripe is projected onto the object surfaces. The first image shows the ideal case where the laser reflection on the surface can be clearly detected for each camera scan line. The second and the third images show the laser reflections on specular surfaces. Notice in these two images multiple peaks (i.e., laser reflections) in the same camera scan lines exist due to mutual reflections. Choosing the peak with the highest intensity value in a scan line – which is the conventional peak detection method – does not guarantee that this is the illuminated point corresponding to the primary reflection of the laser. The conventional approach, therefore, is prone to generate false range measurements in the presence of mutual reflections.

2. Related work

Determining the shape of specular objects has long been a challenging problem in computer vision. Nayar *et al.* [8] proposed an iterative algorithm that recovers the shape and

reflectance properties of surfaces in the presence of mutual reflections. This algorithm is useful for the shape-from-intensity approach to range acquisition; this approach, however, does not produce dense and accurate range maps compared to the optical triangulation methods. Additionally, the proposed algorithm was tested only on Lambertian surfaces of simple geometry.

Clark *et al*. [6] developed a laser scanning system that uses polarization analysis to disambiguate the true illuminated points from those caused by mutual reflections. Their system was tested on shiny aluminum objects with concavities, and false illuminated points were successfully discriminated. However, the system requires special equipment such as linear polarizers, and multiple images need to be captured at each position of the laser. In their experiment, three images were acquired at three different angles of the linear polarizer.

Trucco and Fisher [12] proposed a number of consistency tests for acquiring reliable range images of specular objects. Their range sensor consists of two CCD cameras observing a laser stripe from opposite sides. The consistency tests are based on the fact that the range measurements obtained from the two cameras will be consistent only if the measurements correspond to the true illuminated point. Their method was tested on a polished aluminum block with holes. However, their method does not consider the situation where more than one illuminated point is observed. The consistency tests, therefore, are applied only to the measurements corresponding to a single illuminated point observed per camera scan line. In our experiments, we have noticed that the illuminated points caused by mutual reflections occur very frequently in the vicinity of the true illuminated points, and thus they are seen together along the same camera scan line. Eliminating all points whenever multiple peaks are observed in the same camera scan line may result in too few range measurements.

As an improvement over the conventional methods, we have recently proposed *multi-peak range imaging* [9] – a new range acquisition concept that can also handle surface specularities. False measurements generated by the effects of mutual reflection are eliminated using various constraint tests based on local smoothness, global coordinate consistency and visibility consistency. However, the parameters in the constraint tests were selected manually, and no experimental justification of those selections was provided. Also, a straight-forward implementation of the visibility consistency occasionally caused situations where the true measurements received high inconsistency values. The main contribution of this paper is to resolve these limitations of our previous work.

Due to space limitation, we refer the reader to a good survey paper [2] for the literature concerning the 3D model construction using range images.

3. 3D Modeling Process

Figure 2 shows the flowchart of the 3D reconstruction process. First, several scans from different viewpoints are

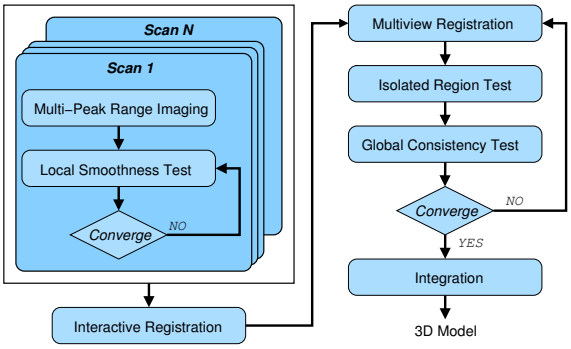

Figure 2. Flowchart of 3D modeling process

performed so that the entire surface of the object can be captured. For each scan, the multi-peak range imaging [9] is carried out to account for the effects of mutual reflections. The local smoothness test is then invoked. This test iteratively eliminates all points that are determined to be locally non-smooth.

After executing the steps described above for each scan, the registration of all the range data acquired from different views is carried out first interactively to yield an approximate registration, and then via a multi-view registration step to fine-tune the registration. The isolated region test is then applied to the registered data followed by the global consistency test which eliminates the points that are determined to be highly inconsistent based on two criteria – namely the coordinate consistency and the visibility consistency – using the information given by all the range data collected from different viewpoints. From now on, we will interchangeably use the term *global tests* as the isolated region and the global consistency test together. If any points were eliminated during the global tests, the multi-view registration step is carried out again on the new data set followed by another application of the global tests. This iteration is continued until no points are eliminated. Finally, the integration is performed on the resulting output. We used the method of Curless and Levoy [7] for the integration.

Here let us define some notations that will be used for the rest of this paper. Let p_i denote the i'th range measurement in a range image. When there is a superscript, for example p_i^j, it denotes the i'th measurement from the range image acquired from the j'th viewpoint. Let $\mathbf{x}(p_i)$ and $\mathbf{n}(p_i)$ denote, respectively, the 3D coordinates and the unit surface normal vector at p_i. The 3D coordinates and the surface normal vector are all with respect to the world frame — the common coordinate system to which all range images are registered. Let $D(a, b)$ denote the closest Euclidean distance between two elements a and b where the elements can be 3D coordinates or line vectors. Finally, let $\Theta(\mathbf{n}(p_i), \mathbf{n}(p_j))$ denote the angle between $\mathbf{n}(p_i)$ and $\mathbf{n}(p_j)$.

3.1. Local smoothness test

It is a legitimate assumption that a range measurement lies on one of the three types of surfaces: smooth surface[1], near crease edge or near jump edge. A crease edge is where surface normals suddenly change directions, and a jump edge is where a spatial discontinuity occurs between adjacent range measurements. A common approach to estimate the local surface property of a range measurement is to fit a planar patch on the neighboring points where the neighboring points typically are those within a small window (e.g., 3×3 or 5×5) centered at the point in question. It has been shown in [5] that the planar patch can be reliably computed not only for those on smooth surfaces but also for those that are located near edges by appropriately selecting neighbors. What we are trying to convey here is that given a true range measurement, whether it lies on a smooth surface or near an edge, we should be able to find a local planar patch that reasonably fits well on the carefully selected neighboring points. If no suitable planar patch can be found, it is likely that the measurement is spurious and should be eliminated. The elimination of such spurious measurements is the object of the local smoothness test.

There are two constraints in the smoothness test. The first constraint requires that each range measurement p_i has the number of valid elements in its fitting window, denoted as $m(p_i)$, greater than a threshold τ_m:

$$m(p_i) > \tau_m \tag{1}$$

Valid data elements in a fitting window are those that have distances to the point of interest less than $b \cdot \rho$ where b is the city-block distance to the point of interest and ρ is the maximum distance allowed between two immediate neighbors. We empirically set ρ to be four times the range sensor resolution. Suppose that there are n valid elements in a fitting window for p_i, including p_i itself. Let us denote those elements as p_e, $e = 1, \ldots, n$, and the center of mass of the elements as $\mathbf{x}_c$. Then, the covariance matrix $\mathbf{C}$ is computed by

$$\mathbf{C} = \sum_{e=1}^{n} [\mathbf{x}(p_e) - \mathbf{x}_c] [\mathbf{x}(p_e) - \mathbf{x}_c]^T$$

The eigenvector corresponding to the smallest eigenvalue of $\mathbf{C}$ represents the normal of the best fitting plane for the elements. Thus, this eigenvector is used as the estimate of the surface normal at p_i, denoted as $\mathbf{n}(p_i)$. The Euclidean distance between an element p_e and the best fitting plane is simply the scalar projection of the vector $\mathbf{x}(p_e) - \mathbf{x}_c$ onto the plane's surface normal $\mathbf{n}(p_i)$. Thus, the fitting error of the elements to the best fitting plane, denoted as $\varepsilon(p_i)$, can be computed by

$$\varepsilon(p_i) = \frac{1}{n} \sum_{e=1}^{n} [\mathbf{x}(p_e) - \mathbf{x}_c] \bullet \mathbf{n}(p_i)$$

[1] Additionally, smooth surfaces may be categorized into 8 different types based on surface curvature sign [3].

The second constraint in the local smoothness test requires that the fitting error of the best fitting plane be less than τ_ε:

$$\varepsilon(p_i) < \tau_\varepsilon \tag{2}$$

All range measurements that do not satisfy either of the two constraints are eliminated. In general, the threshold τ_m must be high enough so that the best fitting plane can be reliably computed, but low enough so that the points near a jump edge will not be eliminated, and the threshold τ_ε must be set in such a way that the points on crease edges will not be eliminated. In some sense, the main task of the local smoothness test is to eliminate only the measurements that no local planar patch is able to fit onto its neighbors. Section 4.2 discusses in more detail how to set the two thresholds τ_m and τ_ε.

3.2. Registration

The registration process in our system consists of two steps: the interactive step that provides an approximate registration and the multi-view registration step which fine-tunes the registration. The interactive step allows a user to look at a set of range images that need to be registered and to click on the corresponding points between the anchor data set and the moving data set. The approximate registration provided by the human interaction serves as the initial registration for the multi-view registration step based on the ICP algorithm [4]. Our multi-view registration is similar to the one proposed by Bergevin *et al.* [1]. Adapting the correspondence criteria presented in [11], our method selects the corresponding points between two data sets as the closest points with the angle between surface normals less than a threshold. The thresholds for selecting the corresponding points are set dynamically in each iteration of the ICP using an approach similar to the one proposed by Zhang [13]. Since our correspondence criteria also includes the angle between surface normals, the angle threshold is also computed dynamically using the same approach as the distance.

Even after multi-view registration, we must anticipate some registration errors, which depend, in general, on the accuracy of the previous registration and the number of remaining false measurements in the data. The registration errors play an important role in the global tests because the tests use the information between all range images, and that information is greatly influenced by how well the range images are registered with one another. It should be mentioned that it is not trivial to compute the registration error for two reasons. First, we do not know which parts between the data sets are overlapping and which parts are distinctive, and second, we do not know which measurements correspond to true surface points and which ones are spurious. Nevertheless, it is commonly accepted in the literature on registration algorithms that the mean distance between the corresponding points or the distance threshold for the correspondence search at the termination of the ICP algorithm be used as the estimate of the registration error. We adopted the latter approach where we use the distance and the angle thresholds at the termination of the ICP algorithm as the

estimates of the registration error. We will denote the registration error of the i'th range image with respect to the distance as λ_D^i, and with respect to the angle as λ_θ^i.

3.3. Isolated region test

The main purpose of the isolated region test is to eliminate all points that are far separated from the true object points. The test involves constructing a 3D volumetric grid that contains the entire data with each voxel having a binary value of 1 if any point exists inside the voxel, and 0 otherwise. A voxel segmentation based on region growing [10] is then performed to cluster the connected voxels. In order to ensure that all true measurements belong to a common connected voxel region, the resolution of the volumetric grid is set as

$$\max\left(\max_i(\lambda_D^i), \rho\right) \tag{3}$$

where λ_D^i is the approximate registration error of i'th range image that was just defined in the last section, and ρ is the maximum distance allowed between two immediate neighbors that was defined in Section 3.1. The isolated region test eliminates all measurements except the ones that belong to the region with the largest size. By the largest size, we mean the largest number of connected voxels in a region.

3.4. Global consistency test

The global consistency test is based on two criteria: the coordinate consistency and the visibility consistency. The coordinate consistency states that the 3D coordinates of true measurements are always consistent among all registered range images that capture the same object surface. On the other hand, the 3D coordinates of false measurements are likely to be inconsistent since the locations where mutual reflections occur depend on the object surface normal relative to the direction of the light source. Assuming that we have range data from N different viewpoints, and that there are a total of M^v measurements in the v'th range image, the coordinate consistency value, denoted as $\mathcal{C}$, is computed by

$$\mathcal{C}(p_i^j) = w(p_i^j) + \sum_{v=1, v\neq j}^{N} \max_{u=1,\ldots,M^v} \left\{\delta_{\mathcal{C}}(p_i^j, p_u^v) \cdot w(p_u^v)\right\} \tag{4}$$

where w is the weight[2] (i.e, confidence value) of a measurement and the test function $\delta_{\mathcal{C}}$ is given by

$$\delta_{\mathcal{C}}(p_i^j, p_u^v) = \begin{cases} 1 & \textbf{if } D(\mathbf{x}(p_i^j), \mathbf{x}(p_u^v)) < \max(\lambda_D^j, \lambda_D^v) \textbf{ and} \\ & \quad \Theta(\mathbf{n}(p_i^j), \mathbf{n}(p_u^v)) < \max(\lambda_\theta^j, \lambda_\theta^v) \\ 0 & \textbf{otherwise} \end{cases}$$

Recall that D is the closest Euclidean distance between two elements and Θ is the angle between surface normals. Also recall that λ_D^i and λ_θ^i are the estimates of the registration error of the i'th range image with respect to the distance and the angle, respectively. Notice in the test function $\delta_{\mathcal{C}}$ that the registration error for each range measurement is incorporated by using them as thresholds. Notice also in Eq. (4) that by taking the maximum value, we are limiting the coordinate consistency value to be contributed only once per range image in case multiple valid points in the same range image exist. In other words, for each measurement p_i^j, the maximum of $N-1$ measurements can contribute to the coordinate consistency value of p_i^j. Since the weight w is normalized to 1 and since the coordinate consistency includes the weight value of its own, the upper bound of the coordinate consistency value is N. This allows us to obtain a more balanced distribution of the coordinate consistency values throughout the data, and more importantly, it allows us to combine the coordinate consistency and the visibility consistency as we shall explain shortly.

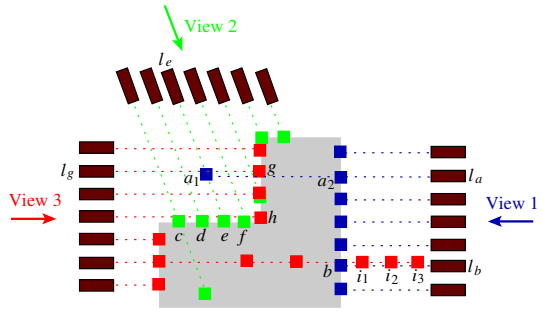

Figure 3. Illustration of visibility consistency

The second criterion of the global consistency test, we call it visibility consistency, is based on the fact that the line space between the sensor and a true measurement is empty, and that the line space beyond a true measurement is invisible to the sensor. Although we can apply this concept for both the projector's line-of-light and the camera's line-of-sight, only the former is considered in this paper; obtaining the visibility consistency for the camera's line-of-sight is computationally much more expensive than that for the projector's line-of-light.

Consider the example depicted in Figure 3 where an object was scanned from three different viewpoints. The dotted lines represent the projector's lines-of-light at the respective sampling positions during the scan. Suppose in the first view, among other detected measurements, points a_1 and a_2 were detected in the same rigel. If a_1 is a true measurement (we do not know yet which one is true), the space at a_2 should have been empty. On the other hand, if a_2 is true, a_1 should not have been visible (or illuminated) by the projector l_a. Thus the measurements a_1 and a_2 are inconsistent with each other with respect to the projector l_a. In fact, the measurements in the same rigel are always inconsistent with each other for there can not be more than one true measurement in the same rigel. The visibility concept applies also for the measurements obtained from different viewpoints. For example, a_1 acquired from the first view and e from the second view are inconsistent with respect to the projector l_e. Similarly, a_1 and g are inconsistent with

[2]Weights are computed as the dot product between the surface normal and the bisection of the camera's line-of-sight and the light projector's line-of-light

respect to the projector l_a and also with respect to l_g.

In order to test the visibility consistency of a range measurement p_i^j, we first need to check for extra measurements in the same rigel since there is always an inconsistency between measurements in the same rigel. We then need to check for measurements from other range images that lie on the line-of-light where p_i^j was sampled. Let the light projector that sampled p_i^j be labeled as l_i^j, and the 3D coordinates of that light projector with respect to the world frame be $\mathbf{x}(l_i^j)$. Then the ray equation of the line-of-light that sampled p_i^j, denoted as $\vec{\mathbf{l}}(p_i^j)$, is given by

$$\vec{\mathbf{l}}(p_i^j) = \mathbf{x}(l_i^j) + \alpha\left(\mathbf{x}(p_i^j) - \mathbf{x}(l_i^j)\right), \qquad \alpha = (0, \infty)$$

The closest distance between $\vec{\mathbf{l}}(p_i^j)$ and a range measurement, say p_u^v, can be computed by

$$D\left(\mathbf{x}(p_u^v), \vec{\mathbf{l}}(p_i^j)\right) = \frac{\left\|\left(\mathbf{x}(p_i^j) - \mathbf{x}(l_i^j)\right) \bigotimes \left(\mathbf{x}(p_u^v) - \mathbf{x}(l_i^j)\right)\right\|}{\left\|\mathbf{x}(p_i^j) - \mathbf{x}(l_i^j)\right\|}$$

where $\bigotimes$ is the cross product and $\|\cdot\|$ the vector magnitude. Obviously, few range measurements will be exactly on a ray, therefore we use the registration error for determining whether a point lies on a line-of-light. That is, we define a range measurement p_u^v to lie on the line-of-light $\vec{\mathbf{l}}(p_i^j)$ if

$$D\left(\mathbf{x}(p_u^v), \vec{\mathbf{l}}(p_i^j)\right) < \max(\lambda_D^j, \lambda_D^v) \tag{5}$$

Once we find the measurements that lie on the line-of-light $\vec{\mathbf{l}}(p_i^j)$, the distance between each of those found measurements and $\mathbf{x}(l_i^j)$ (i.e., the 3D coordinates of the light projector that sampled p_i^j) needs to be computed. If that distance is less than the distance between $\mathbf{x}(p_i^j)$ and $\mathbf{x}(l_i^j)$, there is a visibility inconsistency. For example, in Figure 3, the points a_1 and e are inconsistent because a_1 lies on the line-of-light of e and $D(\mathbf{x}(e), \mathbf{x}(l_e)) > D(\mathbf{x}(a_1), \mathbf{x}(l_e))$. On the other hand, although a_2 and g lie on each other's line-of-light, they are consistent because $D(\mathbf{x}(g), \mathbf{x}(l_g)) < D(\mathbf{x}(a_2), \mathbf{x}(l_g))$ and $D(\mathbf{x}(a_2), \mathbf{x}(l_a)) < D(\mathbf{x}(g), \mathbf{x}(l_a))$. Again we need to take account for the registration error when comparing the distances to the projector. Therefore, we define that there is a visibility inconsistency between p_i^j and p_u^v with respect to the light projector l_i^j if Eq. (5) satisfies and

$$D\left(\mathbf{x}(p_i^j), \mathbf{x}(l_i^j)\right) > D\left(\mathbf{x}(p_u^v), \mathbf{x}(l_i^j)\right) + \max(\lambda_D^j, \lambda_D^v) \tag{6}$$

A straightforward implementation of the visibility consistency can result in situations where true measurements may falsely be determined to be highly inconsistent. For example, the point h in Figure 3 has inconsistency with four measurements c, d, e and f from the second view even though all of them are true measurements. This kind of situation may occur between two orthogonal surfaces where each side of the surfaces is captured by a different range image. For another example, the point b is inconsistent with three false measurements i_1, i_2 and i_3 that happened to be along the line-of-light of b. Similar situations can occur more frequently as the noise increases in the scene and the object shape becomes more complex.

In order to consider all the discussions above, we compute the visibility consistency value for each measurement as follows. Using the same notations as Eq. (4), the visibility consistency value of p_i^j, denoted as $\mathcal{V}(p_i^j)$, is defined as

$$\mathcal{V}(p_i^j) = \sum_{v=1}^{N} \min_{u=1,\ldots,M^v} \left\{\delta_{\mathcal{V}}(p_i^j, p_u^v) w(p_u^v) \left|\mathbf{n}(p_i^j) \bullet \mathbf{n}(p_u^v)\right|\right\}$$

where the test function $\delta_{\mathcal{V}}$ is given by

$$\delta_{\mathcal{V}}(p_i^j, p_u^v) = \begin{cases} -1 & (1)\ \textbf{if}\ p_i^j \text{ and } p_u^v \text{ are in the same rigel } \textbf{or} \\ & (2)\ \textbf{if}\ \text{there is visibility inconsistency} \\ & \qquad \text{between } p_i^j \text{ and } p_u^v \\ 0 & \textbf{otherwise} \end{cases}$$

There are several things that need to be mentioned: First, notice that $\mathcal{V}(p_i^j)$ can only have a zero or a negative value since the test function $\delta_{\mathcal{V}}$ checks only for inconsistency. Second, the absolute value of the dot product between the surface normal vectors is multiplied so that any two measurements whose surface normals are close to orthogonal — such as points c and h in Figure 3 — have little effect to each other. Also, only the smallest consistency value (highest inconsistency) computed for each range image is added to $\mathcal{V}(p_i^j)$ in order to prevent the situation described earlier with the example of the points b, i_1, i_2 and i_3 in Figure 3. By adding only the smallest consistency value for each range image, the lower bound of the visibility consistency value is set to be -N, which enables to obtain a more evenly distributed values of $\mathcal{V}$ throughout the data.

The total global consistency value, denoted as $\mathcal{G}$, is simply the sum of the coordinate consistency and the visibility consistency:

$$\mathcal{G}(p_i^j) = \mathcal{C}(p_i^j) + \mathcal{V}(p_i^j) \tag{7}$$

The main reason for considering the global consistency value $\mathcal{G}$ is that we discovered in our experiment that it is easier to distinguish between the true and the false measurements using $\mathcal{G}$ as opposed to just $\mathcal{C}$ or $\mathcal{V}$. We are able to add $\mathcal{C}$ and $\mathcal{V}$ since the scales of both values are normalized to the number of range images. By simply adding $\mathcal{C}$ and $\mathcal{V}$ without any coefficients, we are assuming they have equal weights. Given the values of $\mathcal{G}$, the global consistency test employs a simple threshold technique in order to eliminate the measurements with low consistency values. There are two constraints in the global consistency test. The first constraint requires that the global consistency value of a measurement be greater than a threshold $\tau_{\mathcal{G}}$:

$$\mathcal{G}(p_i^j) > \tau_{\mathcal{G}} \tag{8}$$

with the threshold set as

$$\tau_{\mathcal{G}} = \min(\mu_{\mathcal{G}} - t \cdot \sigma_{\mathcal{G}}, 0) \tag{9}$$

Figure 4. Objects used in the experiment

where $\mu_{\mathcal{G}}$ and $\sigma_{\mathcal{G}}$ are the mean the standard deviation of $\mathcal{G}$; and t is a positive constant. We shall evaluate the effect of varying t in Section 4.3. Note that by forcing $\tau_{\mathcal{G}} \leq 0$, we are preserving all measurements that do not have any neighbors from other range images as long as there is no visibility inconsistency with other measurements. This property lifts an otherwise restrictive requirement that every part of the object surface must be sampled at least twice from different viewpoints.

The second constraint of the global consistency test also eliminates the measurements with small $\mathcal{G}$ values, but the difference between the first constraint is that it now considers only the range measurements in the same rigel. We define a measurement in a rigel to have a small $\mathcal{G}$ if the value is smaller than the maximum $\mathcal{G}$ in that rigel minus $t \cdot \sigma_{\mathcal{G}}$ where t is the same constant used in Eq. (9). Formally, let $\mathcal{G}_{max}(p_i^j)$ denote the maximum global consistency value among all the range measurements in the rigel to which p_i^j belongs, then the second constraint requires that

$$\mathcal{G}(p_i^j) > \mathcal{G}_{max}(p_i^j) - t \cdot \sigma_{\mathcal{G}} \tag{10}$$

Note that the second constraint only applies to those with more than one measurement in the same rigel. All measurements that do not satisfy either of the constraints, Eqs. (8) and (10), are eliminated in the global consistency test.

4. Experiments

We will now report experimental results on three objects: a bowl, a seashell and an angel figurines, all shown in Figure 4. All three objects have surface materials and shapes that are highly likely to generate mutual reflections.

4.1. Data acquisition

We acquired 5, 18, and 27 range images, respectively, from different viewpoints for the bowl, the seashell and the angel. Then, we painted all the objects so that the surfaces of the objects are ideal for range sensing. Range data of the painted objects were acquired and registered. Let p_{ref} be the set of measurements in all the registered range images of one of the painted objects, and let p_i be a set of measurements from i'th range image acquired from the corresponding original object. Note that none of the elimination tests are applied to these range images yet. Each measurement in p_i is labeled either as true or false by the following procedures: First, p_i is transformed into the coordinate system of the reference data where the transformation matrix is computed beforehand by registering p_i to p_{ref}. In order to obtain an accurate registration, we used the final result of p_i after the convergence of all the constraint tests where the parameters were chosen manually. Once p_i is transformed, each measurement in p_i is labeled as true if there is a point in p_{ref} that has the distance less than the range resolution and the angle between surface normals less than 30 degrees. Otherwise, it is labeled as false. Having all the original data labeled as true of false, we can simply keep record of which of the true or false measurements are eliminated during the constraint tests.

4.2. Analysis of local smoothness test

Figure 5(a) and (b) show the histograms of true measurements (blue line with $\circ$ markers) and false measurements (red line with $\times$ markers) for all three objects as functions of the two parameters used in the local smoothness test. Figure 5(a) shows the histograms as functions of the number of valid elements in the 5×5 fitting window (i.e., the parameter m in Eq. (1)). Notice that the majority of true measurements have all 25 elements, which is the maximum number for a 5×5 window, and few of them have less than 20 elements. On the other hand, there is a considerable amount of false measurements all throughout the range of m from 1 to 20. Since we do not want to eliminate the true measurements in the vicinity of jump edges, a good choice for the threshold τ_m would be from 12 to 20 (i.e., 50% to 80% of maximum number). Figure 5(b) shows the histograms of true and false measurements as functions of the fitting error of the best fitting plane (i.e., the parameter ε in Eq. (2)). Most of the true measurements have the fitting error less than $0.1mm$ whereas the false measurements are more evenly distributed. The graph suggests that a good choice for the threshold τ_ε would be from $0.15mm$ to $0.3mm$ (i.e., about 50% to 100% of sensor resolution[3]).

Figure 5(c) to (e) show the results of the local smoothness test for the bowl, the seashell and the angel, respectively where the parameters are set as $\tau_m = 13$ and $\tau_\varepsilon = 0.2mm$. Approximately three quarters of false measurements were successfully eliminated while maintaining almost all of true measurements.

4.3. Analysis of global tests

The graphs shown in Figure 6(a) to (f) show the number of true and false measurements over the course of the global tests with a varying parameter t of Eqs. (9) and (10). The three graphs from (a) to (c) show the number of true measurements for each of the three objects, and the three graphs from (d) to (f) show the number of false measurements. The odd numbered iterations represent the isolated region test and the even numbered iterations the global consistency test. The initial data was the result of the local

[3]By resolution, we simply mean the distance between two adjacent sampling positions in the scan. For all the range data presented in this paper has the resolution of $0.3mm$.

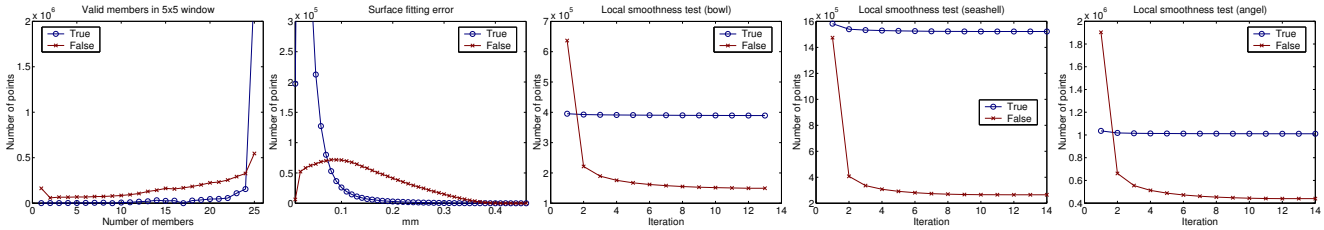

(a) (b) (c) (d) (e)

Figure 5. Analysis of local smoothness test

(a): Histograms of true measurements (blue line with ∘ markers) and false measurements (red line with × markers) as functions of m (the number of valid elements in the fitting window). (b): Histograms as functions of ε (the fitting error of the best fitting plane). (c): Number of true and false measurements for the bowl during the local smoothness test. (d): Number of true and false measurements for the seashell during the local smoothness test. (e): Number of true and false measurements for the angel during the local smoothness test.

smoothness test. Note that the vertical axis limit for each graph is independent, but the lower bounds for the graphs showing the number of false measurements are all zeros. In summary, the graphs suggest that the global tests perform well in the range of $t = 1.0 \sim 2.0$. A smaller t value starts to eliminate too many of true measurements, and a value greater than 2.0 may fail to eliminate a large number of false measurements.

Figure 6(g) visualizes the original data for the bowl where the true measurements are displayed with light blue color and the false measurements with dark red color; Figure 6(h) shows the data after convergence of the local smoothness test; Figure 6(i) the data after convergence of the global tests; and finally Figure 6(j) the reconstructed 3D model. The images in the next row visualize for the seashell and the last row for the angel. For all data, the thresholds were set as $\tau_m = 13$, $\tau_\varepsilon = 0.2mm$ and $t = 2.0$.

5. Conclusion and future work

In this paper, we have presented a 3D reconstruction method that is capable of modeling specular objects. The multi-peak range imaging is used to account for the effects of mutual reflections. A series of constraint tests is applied iteratively to eliminate false measurements generated by mutual reflections. By obtaining the ground truth data, we were able to justify the selection of thresholds in the constraint tests. We showed through our experiments that our approach can generate accurate 3D models of specular objects with complex shapes.

One of the possible future works would be to improve the threshold selections in the constraint tests so that the elimination of false measurements is maximized while preserving all true measurements. For example, the constant t of Eqs. (9) and (10) can be set dynamically in each iteration depending on the current status of the data. Or, instead of using a single t value for the entire data, one can divide the data using regular 3D blocks each of which having different values of t according to the contents of the data within the block.

Another important future work is to extend our methods to other types of optically challenging surfaces such as translucent or highly absorptive surfaces.

References

[1] R. Bergevin, M. Soucy, H Gagnon, and D. Laurendeau. Towards a general multiview registration technique. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(5):540–547, 1996.

[2] F. Bernardini and H. Rushmeier. The 3D model acquisition pipeline. *Computer Graphics Forum*, 21(2):149–172, 2002.

[3] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(2):167–192, 1988.

[4] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[5] C. H. Chen and A. C. Kak. 3D-POLY: A robot vision system for recognizing objects in occluded environments. Technical Report TR-EE 88-48, Purdue University, 1988.

[6] J. Clark, E. Trucco, and L. B. Wolff. Using light polarization in laser scanning. *Image and Vision Computing*, 15(1):107–117, 1997.

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

[8] S. Nayar, K. Ikeuchi, and T. Kanade. Recovering shape in the presence of interreflections. In *IEEE International Conference on Robotics and Automation*, volume 2, pages 1814–1819, April 1991.

[9] J. Park and A. C. Kak. Multi-peak range imaging for accurate 3D reconstruction of specular objects. In *6th Asian Conference on Computer Vision*, 2004.

[10] A. Rosenfeld and A. C. Kak. *Digital Picture Processing*. Academic Press, 1982.

[11] A. Stoddart, S. Lemke, A. Hilton, and T. Penn. Estimating pose uncertainty for surface registration. In *British Machine Vision Conference*, pages 23–32, 1996.

[12] E. Trucco and R. B. Fisher. Acquisition of consistent range data using local calibration. In *IEEE International Conference on Robotics and Automation*, pages 3410–3415, 1994.

[13] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Internation Journal of Computer Vision*, 13(2):119–152, 1994.

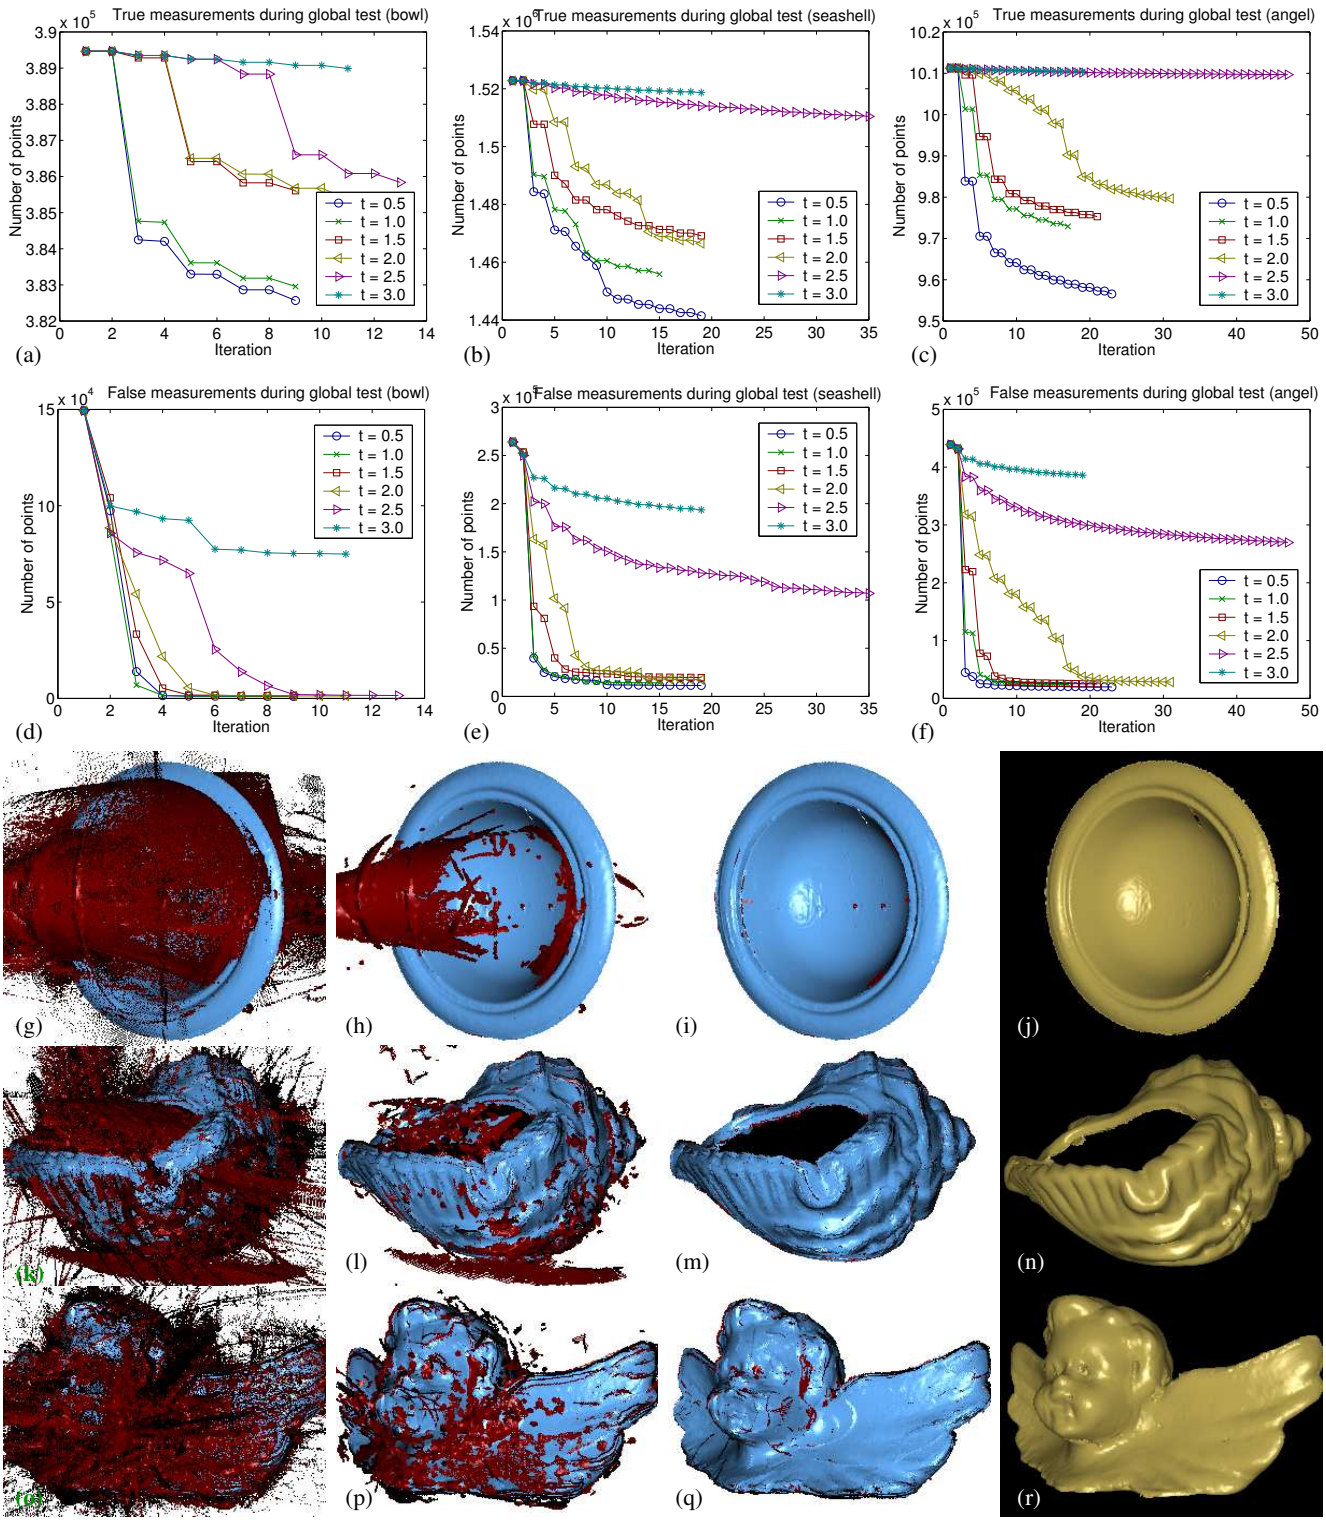

Figure 6.

(a),(b),(c): Number of true measurements for each of the three objects during the course of the global tests with varying parameter t of Eqs. (9) and (10). (d),(e),(f): Number of false measurements during the course of the global tests. (g): The original data for the bowl where the true measurements are displayed with light blue color and the false measurements with dark red color. (h): The bowl after the convergence of the local smoothness test. (i): The bowl after the convergence of the global tests. (j): Final 3D model for the bowl. (k),(l),(m),(n): The same visualization for the seashell. (o),(p),(q),(r): The same visualization for the angel.