

# Recovering Spheres from 3D Point Data

Christoph Witzgall  
National Institute of  
Standards and Technology  
witzgall@nist.gov

Geraldine S. Cheok  
National Institute of  
Standards and Technology  
cheok@nist.gov

Anthony J. Kearsley  
National Institute of  
Standards and Technology  
ajk@nist.gov

## Abstract

The National Institute of Standards and Technology is involved in developing standard protocols for the performance evaluation of 3D imaging systems, which include laser scanners and LADARs (LAsER Detection and Ranging). A LADAR is an optical device that typically yields voluminous 3D “point clouds” by scanning scenes. In many applications, a model of an object which is present in the scene has been specified, and the task amounts to recovering this object from scan data. Specifically, the recovery of spheres from point clouds will be addressed, aiming at estimating the location of their centers and, if not known beforehand, their radii. This information can be used, for instance, to “register” LADAR data to a specified coordinate frame. Two experiments recovering spheres based on best-fitting data points are reported. Sphere fitting based on orthogonal least squares is compared to a novel approach, minimizing instead the squares of range errors incurred in the direction of the scan.

## 1. Introduction

The past decade has witnessed the rapid growth of the technology for 3D imaging systems (e.g., laser scanners and LADARs - LAsER Detection and Ranging) and the applications for these systems. LADAR determines 3D coordinates of points on an object by sending out a signal and analyzing its reflection back to the instrument as shown in Figure 1. A LADAR is mainly a scanning device in that it sends out a multitude of signals in a very short time, resulting in point clouds containing typically millions of data points. Applications include monitoring of construction sites, the development of as-built models of existing structures and conditions, mapping, visualization of “hidden” objects, guiding of unmanned vehicles, just to mention a few.

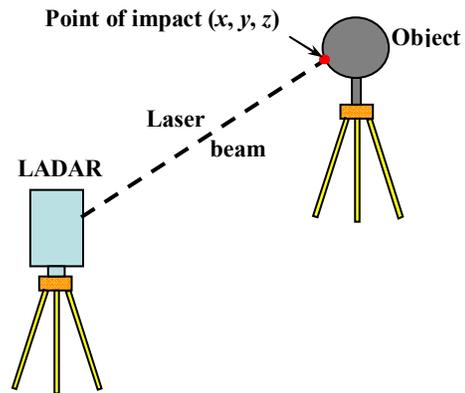


Fig. 1. Schematic of the operation of a LADAR scanner.

The metrology of LADARs is a current research issue at the National Institute of Standards and Technology (NIST). This work also supports the development of standard test protocols for the performance evaluation of LADARs. Figure 2 shows the indoor, artifact-based LADAR facility at NIST.

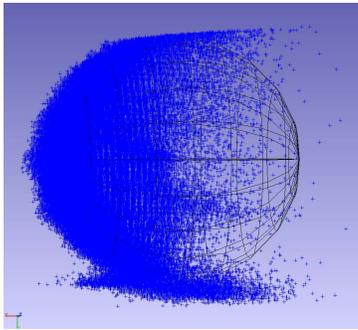


Fig. 2. Indoor, artifact-based 3D Imaging Facility at NIST.<sup>1</sup>

Figure 3 presents a LADAR visualization of a rubble pile. The casual reader may see this picture as just what a photographic camera might produce. But this would miss the point that, once the point cloud has been captured in 3D, it can be displayed as seen from different viewpoints. For instance, the point cloud generated off a sphere is displayed in Figure 4 as seen sideways, that is, perpendicular to the direction of the scan.



**Fig. 3. LADAR scan of a rubble pile.**



**Fig. 4. Point Cloud and fitted sphere.**

The point cloud in Figure 4 consists of actual data points. The superimposed picture of the sphere does not depict the actual sphere but a “fitted” sphere, a sphere that in some sense best represents the point cloud, a sphere that was calculated using a fitting algorithm. The fitted sphere is expected to capture the radius and location of the actual sphere vis-à-vis the point cloud. Fitting the sphere thus provides a method for recovering the actual sphere from which the point cloud had been generated.

## 1.1 The Fitting Paradigm

Fitting requires the definition of a “gauge function,” which expresses a measure of deviation of the point cloud as a whole from a putative geometric model such as a sphere. The parameters of the model governing its shape and size are the arguments of the gauge function. They are the variables of an optimization problem: to minimize the deviation of data points from an object model.

Gauge functions are typically defined in two steps. First, a measure of the individual deviation

$$\Delta_i = \text{deviation}((x_i, y_i, z_i), \text{model})$$

of each data point from the putative model is defined. Then a norm  $\|*\|$  is selected and applied to the vector of all deviations  $\Delta_i$ , in order to arrive at a single gauge value for an overall measure of deviation,

$$\Delta = \|(\Delta_1, \Delta_2, \dots, \Delta_i)\|$$

The quantity  $\Delta$  represents the desired gauge function, which depends on the model parameters as its arguments, since every one of the individual deviations  $\Delta_i$  depends on those parameters.

The following are the most commonly selected norms:

- the Chebychev norm  $(L_\infty)$ :  $\max_i |\Delta_i|$
- the least-squares norm  $L_2$ :  $\sqrt{\sum_i \Delta_i^2}$
- the sum of absolute values norm  $L_1$ :  $\sum_i |\Delta_i|$

The most commonly used gauge function is based on orthogonal distances of data points  $(x_i, y_i, z_i)$ , from the sphere (see Fig. 5). These distances are then combined by summing their squares,

$$\sum_i \left[ \sqrt{(x_0 - x_i)^2 + (y_0 - y_i)^2 + (z_0 - z_i)^2} - r \right]^2,$$

where  $(x_0, y_0, z_0)$  and  $r$  denote the center and radius of the sphere, respectively. Fitting procedures which minimize the above gauge function are commonly referred to as “orthogonal least squares” or “geometric sphere fitting” [1, 9].

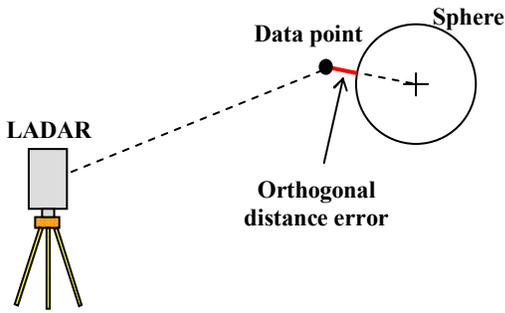


Fig. 5. Orthogonal distance error.

The use of the Chebychev norm has been explored by [2, 3, 10]. Applications are in the quality control for manufactured parts to determine whether they meet desired tolerance criteria.

In our work on spheres, the desired end result consists of the “center coordinates”  $(x_0^*, y_0^*, z_0^*)$  of a fitted sphere, and perhaps its “radius”  $r^*$ . Indeed, when fitting spheres, two different tasks may be encountered:

- fitting a sphere with its radius “free” to be determined;
- fitting a sphere with a specified “fixed” radius.

## 2. Locating I-Beams

The following demonstration experiment [5] was designed to demonstrate the feasibility of automated placing and picking of an I-beam by the computer-guided crane developed at NIST. The I-beam, residing on the laboratory floor, was scanned for location and orientation with the data in the LADAR’s coordinate system.

This instrument’s coordinate system then had to be related to the coordinate system of the crane, a process generally called “registration”. To this end, three “target” spheres, “A”, “B”, “C”, were placed in the vicinity of the I-beam (see Fig. 6). The centers of the spheres were predetermined in the coordinate system of the crane. The LADAR scan covered these spheres along with the I-beam, and the fitting process yielded center coordinates in the instrument system. Thus, there were three target locations, each of which with coordinates known in both systems. The “Procrustes” algorithm [7] was employed, which combines a translation with a rotation in order to transform one system into the other, matching the coordinates at each target point as well as possible.

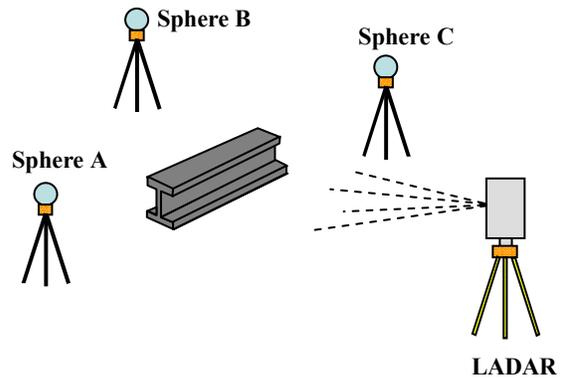


Fig. 6. Determining location and orientation of an I-beam using LADAR.

It is clear that the accuracy of the fitting algorithm as applied to the target spheres is the key to a correct registration. Thus, the fixed radius option for sphere fitting was used for the actual demonstration. The radii of the target spheres (76.2 mm, 3 in) were known, and it could be determined how well the radii were reproduced if the free radius option of the sphere fitting algorithms were to be used. The results were disconcerting.

The radii were generally underestimated. In particular, applying a commercial fitting package to target sphere “C” yielded for  $n = 90$  the average radius,  $r_{average}$ ,

$$r_{average} = 69 \text{ mm} (\text{std. dev.} = 3 \text{ mm}), r_{actual} = 76 \text{ mm},$$

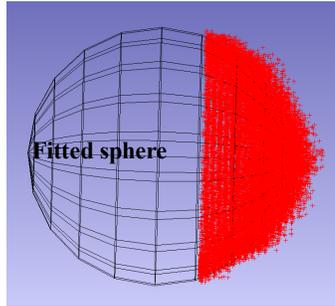
a discrepancy of about 10 %. An effort was therefore started to determine the reasons for such errors.

At first, suspicion centered on the quality of the target spheres which were made of polystyrene. It was thought that this material might permit some penetration of the sphere by the laser beam, or perhaps the dimensions of the spheres were not quite accurate. The experiment described in the following section was therefore conducted.

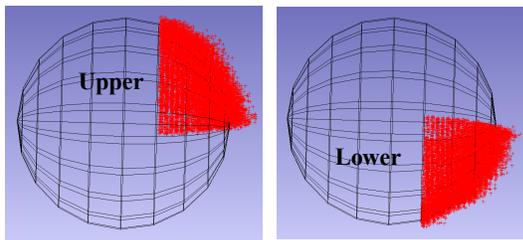
## 3. Fitting an Aluminum Sphere

In this experiment, a data set was collected, and reduced to avoid boundary effects, off an aluminum sphere machined to a radius of 101.6 mm (4 in). This data set is displayed in Figure 7 together with two subsets, an upper and a lower subset into which the full set had been split as shown in Figure 8. The

results of applying the commercial fitting package to these three data sets are displayed in Tables 1 and 2.



**Fig. 7. Full “hemispherical” data set from aluminum sphere.**



**Fig. 8. Upper and lower portions of the hemispherical data set.**

**Table 1. Orthogonal fitting; variable radius.**

	$x$	$y$	$z$	$r$
Full	-6254.99	-196.51	-78.85	98.41
Upper	-6258.27	-196.37	-83.02	102.36
Lower	-6258.61	-196.82	-72.61	103.66

**Table 2. Orthogonal fitting; fixed radius.**

	$x$	$y$	$z$	$r$
Full	-6259.19	-196.58	-78.87	101.6
Upper	-6257.52	-196.36	-82.55	101.6
Lower	-6256.59	-196.77	-73.98	101.6

The first observation concerns the result of fitting the full data set with the free radius option. As in the demonstration reported in Section 2, the radius was still underestimated:

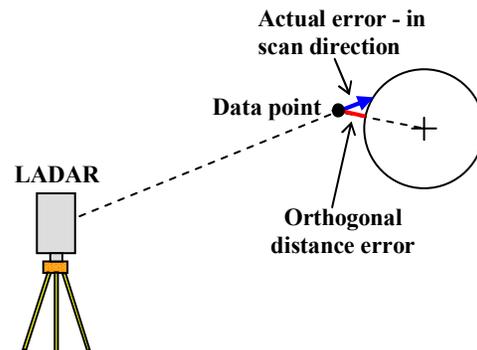
$$r_{computed} = 98.41 \text{ mm}, r_{actual} = 101.6 \text{ mm},$$

but then it was overestimated for both the upper and the lower portion of the full data set. The next observation concerned the high level of sensitivity in the  $z$ -coordinate, which represents vertical elevation. Note that the same sensitivity in the  $z$ -coordinate showed up when the known radius of 101.6 mm had been kept fixed. Such variations are at odds with the fact that regions on the sphere are equivalent. Indeed, the upper and the lower data set occupy essentially symmetric positions on the sphere. Yet there is a substantial difference in fitting results.

The upper and the lower subsets are, however, in a different position vis-à-vis the LADAR instrument. The angles of incidence certainly differ for these two subsets. This forces the conclusion that the instrument position has to be taken into account when fitting. A method for this will be presented in Section 4.

#### 4. Fitting in Scan Direction

In what follows, we assume that the data points are acquired in scan direction, i.e., the data point  $(x_i, y_i, z_i)$  and its intended impact point on a sphere are found on a ray, referred to as “scan ray,” emanating from the origin  $(0, 0, 0)$ , so that the intended target point is at the intersection of the scan ray with the sphere. The distance between the data point and that intersection will be considered the actual error - see Fig. 9. Of course, the scan ray may not intersect the sphere, in which case we declare a “miss;” otherwise, a “hit.” In reality, one might argue, every data point would arise from a hit, as otherwise it would not have gotten into the data set. But as long as a sphere is not in the best-fit position, misses are to be expected.



**Fig. 9. Actual error in scan direction.**

These considerations suggest a gauge function based on the range errors in scan direction, in particular, combining them using the square root of the sum of squares of the respective range errors. This

approach, however, fails to provide a suitable basis for a gauge function as it depends solely on data points scoring hits. Indeed, any sphere position that avoids hits altogether would yield a minimum of zero. For this reason, an appended gauge function is proposed which includes also points which produce misses, and defining their deviations as the orthogonal distance to the sphere. This expenditure will be resorted to until the set of hits stabilizes. The final result will be based on the stabilized set of hits only.

This and similar procedures are still subject to experimentation and adjustments. Our first efforts, however, were clearly successful as attested by the results given in Table 3. They indicate that the abnormalities reported in Section 2 appear indeed to be caused primarily by “modeling error,” namely the choice of an unsuitable gauge function for the fitting procedure.

**Table 3. Fitting in scan direction; variable radius.**

	$x$	$y$	$z$	$r$
Full	-6258.98	-198.07	-79.18	101.29
Upper	-6259.06	-198.15	-78.90	101.22
Lower	-6259.38	-198.01	-79.12	101.60

Unfortunately, the appended gauge function is not continuous because a perturbation of the sphere location may cause a hit to become a miss, or vice versa, and thus switch to a different definition of deviation. Fortunately, those discontinuities are associated with a small number of points and thus tend to have only a noise-like effect on the overall gauge function. Nevertheless, an optimizer is needed which does not require differentiability and also tolerates small discontinuities. Recall also that the final optimization will be with respect to the un-appended gauge function. The method that turned out to be successful was a method based on recent research in the optimization of noisy functions. Loosely speaking, this method is designed for minimizing functions with many local but non-global minima. The algorithm works well as long as these minima are “shallow,” i.e., they act as perturbations of an overall smooth and well-behaved function.

This method also permits constraining the minimization to avoid, for instance, values of variables that define spheres that would infringe upon the instrument location.

The algorithm proceeds in two stages. In order to form a centered finite difference approximation to the gradient of the gauge function, a finite-difference step-length is calculated. Initially this finite-difference step-length is taken to be quite large and is gradually decreased as iterates draw closer to the solution [4]. When the gauge function becomes sufficiently small,

this finite difference step-length is defined to be the value that maximizes the expected accuracy in the gradient of the gauge function given an estimate of the error in the gauge function incurred by each of the “misses”. The algorithm progresses towards the solution until convergence criteria are met and the finite difference step-length is decreased below the square root of machine precision or to a size to be sufficiently below a threshold defined by the estimate of noise in the gauge function generated by misses. Each iteration of the algorithm is based on a minimization a quasi-Newton (in this case, BFGS) method. With the gradient calculation in hand any gradient-based method could have been employed (e.g., conjugate gradient). As a final step, once convergence the gradient convergence criteria has been met, a direct search method coordinate search is employed [8] to guarantee as much accuracy in the solution as possible. This coordinate search method requires no gradient calculation or approximation, and has been applied successfully in the past on difficult non-differentiable constrained optimization problems (e.g., [6]).

## 5. Concluding Remarks

The experimental evidence presented here, although limited so far, indicates that for certain applications it may not be enough to base fitting methods on the coordinates of points alone, but that it may be necessary to take into account the directions in which those points had been acquired. This may require a general revision of currently used approaches to approximation.

## 6. Footnotes

1 Certain products are shown in this photograph. In no case does this imply recommendation or endorsement by NIST, nor does it imply that the products are necessarily the best available for the purpose.

## References

1. W. Gander, G.H. Golub, and R. Strebler. Least-squares fitting of circles and ellipses. *BIT*, 34:558–578, 1994.
2. J. Garcia-Lopez, P. A. Ramos, and J. Snoeyink. Fitting a set of points by a circle. *Discrete and Computational Geometry*, 20:389–402, 1998.
3. S.I. Gass, C. Witzgall, and H.H. Harary. Fitting circles and spheres to coordinate measuring machine data.

International Journal of Flexible Manufacturing Systems, 10:5–25, 1998.

4. P. Gilmore and C.T. Kelley. An implicit filtering algorithm for optimization of functions with many local minima. *SIAM Journal on Optimization*, 5:269–285, 1995.
5. D.E. Gilsinn, C. Witzgall, G.S. Cheok, and A. Lytle. Construction object identification from LADAR scans: an experimental study using I-beams. NISTIR 7286, National Institute of Standards and Technology, Gaithersburg, MD, 2005.
6. R. Glowinski and A.J. Kearsley. On the simulation and control of some friction constrained motions. *SIAM Journal on Optimization*, 5:681–694, 1995.
7. E. Golub and C.F. Van Loan. *Matrix Computation*. The Johns Hopkins University Press, Baltimore, MD, 1996.
8. R. M. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9:1082–1099, 1999.
9. C.M. Shakarji. Least-squares fitting algorithms of the NIST algorithmic testing system. *Journal of Research of the National Institute of Standards and Technology*, 103:633–640, 1998.
10. C.K. Yap. Exact computational geometry and tolerancing metrology. *Snapshots on Computational and Discrete Geometry*, D. Avis and P. Bose, eds., McGill School of Computer Science, 3:34–48, 1994.