

Experiences with Point Cloud Registration¹

by

Christoph Witzgall² and Geraldine S. Cheok³

ABSTRACT: The development of LADAR (laser distance and ranging) technology to acquire 3D spatial data made it possible to create 3D models of complex objects. Because an unobstructed line-of-sight is required to capture a point on an object, an individual LADAR scan may acquire only a partial 3D image, and several scans from different vantage points are needed for complete coverage of the object. As a result there is a need for software which registers various scans to a common coordinate frame. NIST is investigating direct optimization as an approach to numerically registering 3D LADAR data without utilizing fiducial points or matching features. The primary capability is to register a point cloud to a triangulated surface - a “TIN” surface. If a point cloud is to be registered against another point cloud, then the first point cloud is meshed in order to create a triangulated surface against which to register the second point cloud. The direct optimization approach to registration depends on the choice of the measure-of-fit to quantify the extent to which the point cloud differs from the surface in areas of overlap. Two such measures-of-fit have been implemented. Data for an experimental evaluation were collected by scanning a box, and registration accuracy was gauged based on comparisons of the volume and height to known values.

KEYWORDS: LADAR; measures-of-fit; point cloud; registration; TIN; triangular mesh.

1. INTRODUCTION

In recent years, the National Institute of Standards and Technology (NIST) has investigated metrological aspects of LADAR (Laser Distance and Ranging) scanning, addressing both hardware and software issues [2,4,5]. Hardware calibration issues include statistics for direct range measurements, their dependence on target color, distance, and angle of incidence. Corresponding experiments will be the topic of a forthcoming report [2]. In this work, the focus is on software issues, in particular, the registration of LADAR scans taken of the same scene from different vantage points. Statistical experiments aimed at assessing triangular meshing for surface modeling are described in another forthcoming report [5].

The data for the software experiments consist of four LADAR scans taken indoors of a

$$0.914 \text{ m} \times 1.219 \text{ m} \times 1.524 \text{ m} \\ (3 \text{ ft} \times 4 \text{ ft} \times 5 \text{ ft})$$

wooden box. The task is to create a triangular mesh representing the box and surrounding floor. The accuracy of that representation is gauged by calculating from it the volume and height of the box. The accuracy of the box dimensions is $\pm 1.58 \text{ cm}$ (1/16 in) and assuming worst case, the volume error is $\pm 0.4 \%$ [4].

The registration methods considered here match a point cloud against a triangulated elevated surface. The methods are conceived as optimization problems, which apply rigid transformations to a point cloud in order to minimize a chosen

¹ Official contribution of the National Institute of Standards and Technology (NIST); not subject to copyright in the United States.

² NIST, ITL, Mathematics and Computational Sciences Division, Mail Stop 8910, Gaithersburg, MD 20899-8910, witzgall@nist.gov

³ NIST, BFRL, Construction Metrology and Automation Group, Mail Stop 8611, Gaithersburg, MD 20899-8611, cheok@nist.gov

“measure-of-fit”, which quantifies the divergence of the point cloud from the surface. Several distinct measures-of-fit will be identified, and the performance of two of them will be examined using the box data. Evaluations of additional options are in progress.

Direct minimization of a selected measure-of-fit differs from the more commonly used ICP (Iterative Closest Point) approach [1,3]. Direct minimization, however, is suited to the purpose of examining the effects of the choice of measures-of-fit.

In a first series of experiments, each of the four scans is registered individually to an exact triangulated model of the box and surrounding floor space. The registered four point clouds are combined into a single point cloud, which is then “cleaned” and meshed to create a surface model of the box. In a second series of experiments, the recreation of the box is attempted by pair-wise registration of the point clouds against themselves without benefit of an accurate reference. The four point clouds are examined in clockwise rotation. A triangulated surface is determined by cleaning and meshing the first point cloud. The second cloud is then registered against the surface model of the first cloud. Once that common frame has been established, the two point clouds are combined, cleaned and meshed to generate a common surface representation of the first two point clouds, against which to register the third point cloud. The process is then repeated by registering the fourth point cloud against the surface representation of the other three point clouds. Finally, all four point clouds -- now registered to the frame of the first -- are combined, cleaned, and meshed to produce a model of the box.

Earlier experiments have been conducted along the same lines [4]. Here we report the outcomes of fully automated procedures.

2. MEASURES-OF-FIT

There are two distinct aspects to distance based measures-of-fit to be identified in this report:

- The definition of generic point-to-surface distance
- The selection of a “norm” by which to distill a single number from many instances of such distances

Commonly considered norms are the

MAX, RMS, and ASD norms.

The MAX norm assigns the maximum absolute value or “size”. The RMS (“root-mean-square”) norm averages the sum of squares before taking the square root. The ASD (“average-size-deviation”) represents the average of the absolute values. Norms typically provide error estimates in the form of deviations from zero. The MAX norm will not be considered here, since it is too dependent on “outliers”. Each of the other two norms may be combined with a generic point-to-surface distance measure to arrive at a measure-of-fit.

Three generic distance measures are considered:

- Vertical distance
- Euclidean distance
- Ray-directed distance

A surface is “elevated” or “2.5 D” if it has a unique projection into the x,y -footprint plane. The term “TINsurface” (TIN=triangulated irregular network) is frequently used for triangulated surfaces that are also elevated, and only such surfaces are considered in this report.

“Vertical distance” is defined as the absolute value of the “residual” $r_i = z_i - \hat{z}_i$ of a point $p_i = (x_i, y_i, z_i)$, where \hat{z}_i denotes the elevation of that surface point $\hat{p}_i = (x_i, y_i, \hat{z}_i)$ which has the same footprint (x_i, y_i) as the data point p_i . The point \hat{p}_i is uniquely determined, provided the surface is elevated. The point \hat{p}_i , however, does not exist if the footprint (x_i, y_i) of the data point lies outside the footprint region of the surface. This fact gives rise to a natural concept of “overlap”. Indeed, measures-of-fit should apply

only to areas of overlap between the point cloud and the surface.

“Euclidean distance” is the conventional distance between point $p_i = (x_i, y_i, z_i)$ and the surface, that is, the smallest distance between p_i and any point on the surface. This distance is always defined. It is therefore necessary to exclude non-overlap areas. Roughly speaking, “Ray-directed distance” is defined as the distance measured in the direction of the laser beam. The overlap area is determined by those rays which meet both point cloud and surface.

In this report, only vertical distances are considered together with the RMS and ASD norms, respectively. In other words, the following measures of fit are considered:

- Vertical distance with the ASD norm
- Vertical distance with the RMS norm

Vertical distance is a natural choice for 2.5D, that is, TIN surfaces such as terrain representations.

3. OPTIMIZATION PROCEDURE

The intent of optimization-based registration methods is to identify a rigid transformation which -- when applied to a given point cloud -- repositions the point cloud so as to minimize deviation from the given surface as quantified by a measure-of-fit. Rigid transformations may be characterized by a translation with parameters,

$$\Delta x, \Delta y, \Delta z,$$

combined with three rotations,

$$\Delta \varphi, \Delta \varepsilon, \Delta \theta,$$

the yaw, roll, and pitch, respectively [4]. The value of the measure-of-fit after transformation may be therefore considered as a function, F , of the above six transformation parameters, and this function is to be minimized.

In the case of the measures-of-fit based on vertical distance, this minimization met with several difficulties. First, in the neighborhood of an optimal parameter choice, the function F assumes

the shape of a plateau with many local minima in close proximity of each other, and the quality of the registration tends to be sensitive to the choice of one of those local minima. Moreover, a global minimum is typically a disastrous choice. The automated procedure used for this work searches among neighboring local minima within a fixed radius and terminates if no improved measure-of-fit can be found. This procedure is, of course, very sensitive to the choice of starting point: If the starting point fails to be reasonably close to an acceptable “solution”, the procedure may, in fact, lead to progressively worse registrations, all the while “improving” the measure-of-fit. The challenge is to find measures-of-fit which track the quality of registration. Those considerations lead to the following optimization process.

Having evaluated a particular set of parameters $(\Delta x^0, \Delta y^0, \Delta z^0, \Delta \varphi^0, \Delta \varepsilon^0, \Delta \theta^0)$, the translation $(\Delta x, \Delta y, \Delta z)$ is optimized, keeping the angle parameters fixed. Here, the two planar parameters $(\Delta x, \Delta y)$ are considered first. This planar optimization procedure will be described below in more detail. Once that optimization step is completed, the optimal vertical translation Δz may be determined in closed form by subtracting from the previous value either the mean or the median of all residuals, depending on whether the measure-of-fit is RMS or ASD based, respectively. If the new vertical translation parameter Δz changes from its previous value, then the planar optimization procedure resumes. Otherwise, the translation parameters $(\Delta x, \Delta y, \Delta z)$, are considered optimized, and the process moves to the optimization of the angle parameters.

In order to avoid minima that correspond to unacceptable registration results, the optimization of the planar parameters $(\Delta x, \Delta y)$ follows a “limited-horizon” search principle: Consider a circle of radius r around a current “point” $(\Delta x^0, \Delta y^0)$. Search for a “better” point $(\Delta x, \Delta y)$ within this circle only. If found, it becomes the center of another search within a radius of r around it. If no such improvement is found, then the planar search terminates and, as described above, the optimal vertical parameter Δz is determined. The procedure consistently decreases with measure-of-fit function F , and terminates with a particular “local” minimum.

For the experiments reported here a search radius, $r = 5$ cm, was used. If r is too large, it may lead the search astray. If r is too small, the search may terminate before reaching a “good” local minimum. As for most registration methods, the success of the search depends heavily on the choice of the initial parameters (“warm start” vs. “cold start”).

There are, of course, many approaches to optimizing a function in a given region. Most advanced optimization algorithms for minimizing a function $f(\xi_1, \xi_2, \dots, \xi_n)$, however, stipulate that a point $(\xi_1^0, \xi_2^0, \dots, \xi_n^0)$ is a local minimum if perturbing single variables ξ_i^0 does not lead to lower values of f . This is not true for functions that are not continuously differentiable such as the measure-of-fit function F .

For this reason, we adopted a search based simply on sampling the search area. With $(\Delta x^0, \Delta y^0)$ as the center point, concentric sets of 20 trial points at equal angle increments of 22.5° are considered for 15 radii ranging from 0.02 cm to 5.0 cm. Generally, the search starts with radius 1.0 cm and if necessary, moves to successively smaller ones of the proposed radii. If reduction to the smallest radii fails to yield an improvement, then the search continues with increasing radii up to the search radius of 5 cm. As F is calculated for up to $15 \times 20 = 300$ parameter settings for each planar search step, the search procedure is computationally expensive and in this form, too slow for real-time purposes.

It was assumed, that the optimization of the angle parameters does not need to be similarly restricted except for reasons of efficiency. The planar rotation $\Delta\phi$ is varied in single steps of 0.05° through a specified bracket and, for each value, the full translation $(\Delta x, \Delta y, \Delta z)$ is re-optimized. The selection of the roll and pitch parameters $\Delta\epsilon, \Delta\theta$ has as yet not been automated, because previous experiments [4] suggested that -- for reasons as yet unclear -- minimizing those parameters actually worsened the quality of the registration.

Given a TINsurface, each specification of a horizontal “cut plane”, at elevation z , defines both

a “cut volume” and a “fill volume”. The cut volume is bounded below by the cut plane and above by the surface as far as it extends above the cut plane. The fill volume is bounded above by the cut plane and below by the surface.

In the experiments to be described, evaluating a sequence of cut volumes was used for determining volume and height of the box once its surface and that of the surrounding floor has been modeled. More precisely, the idea is to create a sequence of cut volumes, starting with cut planes below the floor level, as indicated by a fill volume of 0, and proceeding by increasing the cut level z , in equal increments of 0.2 cm, until the cut plane clears the top of the box, as indicated by a cut volume of 0. As long as the cut plane remains below the surface, the cut volume will decrease linearly, the slope given by the area of the footprint region of the surface. As the cut plane starts to intersect the surface, a transition occurs to the regime in which the cut plane only intersects the box, and where the decrease is approximately linear with the box footprint as its slope. A similar transition is observed as the cut plane moves beyond the top of the box. The first transition marks the floor level. The second transition marks the top level of the box. The two transition elevations show up as “spikes” in the sequence of second differences [4]. The locations of these spikes define the elevations for floor and box top, respectively. Thus

$$\begin{aligned} \text{box height} &= \text{box top elevation} - \text{floor elevation} \\ \text{box vol.} &= \text{cut vol. @ floor} - \text{cut vol. @ box top} \end{aligned}$$

4. EXPERIMENTS

As mentioned previously, four scans of a box were obtained. These scans were obtained with the scanner located “in front” of each box corner, i.e., in each scan, the box top and two sides of the box were visible. The scans are clockwise labeled C, D, E, F. The four scans were visually transformed so that they were in rough alignment to each other, Figure 1. This was necessary as the registration program assumes small transformations – this is true for most registration programs.

Two series of experiments were conducted: 1) registration against an exact model of the box and

2) registration of a point cloud to another point cloud. For the second series, point cloud C was arbitrarily selected as the reference point cloud. Parameters for both series include: 1) starting point of registration, cold vs. warm start and 2) measure-of-fit, ASD or RMS. For the starting point of registration, a cold start was a start at the origin where $(\Delta x, \Delta y, \Delta z) = (0, 0, 0)$ and a warm start was a start off the origin at some point $(\Delta x, \Delta y, \Delta z)$ close to the expected final registration point [4]. For the cold start, the ϕ value ranged from -2.5° to 2.5° , and the increment size was 0.05° . For the warm start, the range of the ϕ value was 1° starting close to the final value [4]. The values for ϵ and θ were set to zero and not changed.

The results of the automated registration are given in Tables 1 and 2. Summary of results:

- ASD vs. RMS: ASD yields a better registration in terms of volume and box height
- Registration to box: ASD volume error on the order of 2 % (cold and warm starts), RMS volume error of 12 % (warm start) and 137 % (cold start) – suspect registration got stuck on local minima.
- Point cloud to point cloud registration: ASD volume error on the order of 10% (cold and warm starts), RMS volume error 13 % (warm start) and 36 % (cold start).
- Warm vs. cold start: For RMS, significant reduction of volume error for warm start. For ASD, no significant difference in volume error for warm vs. cold start.

Figures 2 (registration to box) and 3 (point cloud to point cloud registration) show the final registration of the four scans. The reason for the misalignment of the last scan in Figure 3 is being examined.

5. CONCLUSIONS

The experiments reported in this paper clearly show the importance of selecting the correct measure-of-fit for optimizing the registration. In these experiments which were based on vertical distance, the ASD norm produced a better registration than the RMS norm in terms of box volume and box height. Future work will include using other measures-of-fit such as the “ray-directed” distance.

6. REFERENCES

1. Besl, P., and McKay, N., “A Method for Registration of 3-D Shapes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 14, No. 2, 1992, pp. 239-256.
2. Cheok, G. S. and Leigh, S., Calibration of a Laser Scanner, in preparation.
3. Rusinkiewicz, S. and Levoy, M., “Efficient Variants of the ICP Algorithm,” in *Proc., International Conference on 3D Digital Imaging and Modeling (3DIM2001)*, IEEE Computer Society, 28 May – 1 June, 2001, pp. 145-152.
4. Witzgall, C. and Cheok, G. S., Registering 3D Point Clouds: An Experimental Evaluation, *NISTIR 6743*, NIST, Gaithersburg, MD, May, 2001, 41 pp.
5. Witzgall, C. and Cheok, G. S., Triangular Meshing for Volume Determination, in preparation.

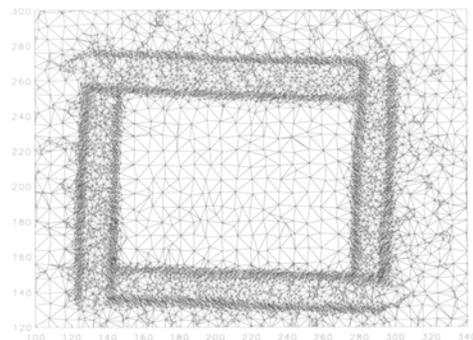


Figure 1. Triangulation of Footprint of Roughly Aligned Scans - Starting Point for Registration.

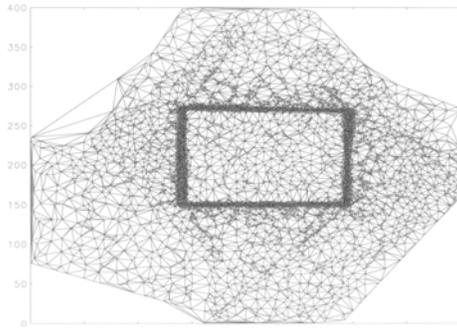


Figure 2. Registration to Box – Final Registration.

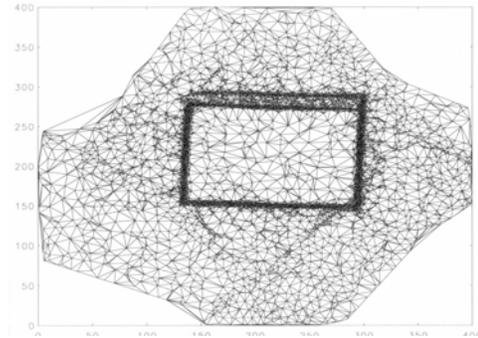


Figure 3. Point Cloud to Point Cloud Registration – Final Registration.

Table 1. Registration of Point Cloud to Exact Box Model.

	Norm Used	Final Registration Values						ASD (cm)	RMS (cm)	Volume (m ³)	Volume Error (%)	Box Height (cm)	Box Ht. Error (%)	
		Δx (cm)	Δy (cm)	Δz (cm)	$\Delta\phi$ (°)	$\Delta\epsilon$ (°)	$\Delta\theta$ (°)							
Start at $(\Delta x, \Delta y, \Delta z) = (0, 0, 0)$ (Cold Start)	ASD	C	6.550	-8.790	-1.000	2.45	0	0	19.101	33.877	1.727	1.63	92.0	0.61
		D	6.605	8.069	0.100	2.35	0	0	25.206	38.890				
		E	14.052	0.443	-3.900	-0.45	0	0	24.910	38.898				
		F	16.984	13.718	-1.900	1.20	0	0	24.089	38.111				
	RMS	C	0.370	-6.318	15.072	2.00	0	0	24.071	30.557	4.026	136.97	92.6	1.27
		D	1.956	2.566	-27.112	2.50	0	0	27.395	31.056				
		E	18.636	-13.208	18.423	2.30	0	0	27.565	33.550				
		F	19.047	12.148	-26.209	2.40	0	0	27.152	30.860				
Start off origin (warm start)	ASD	C	7.443	-9.829	-0.800	2.80	0	0	19.044	33.770	1.671	-1.65	92.8	1.49
		D	6.048	8.564	0.100	2.20	0	0	25.211	38.895				
		E	20.246	-12.027	-1.900	2.25	0	0	25.005	38.838				
		F	25.812	6.232	-0.300	2.95	0	0	24.349	38.335				
	RMS	C	0.372	-6.308	15.072	2.00	0	0	24.071	30.557	1.906	12.20	97.5	6.63
		D	8.936	8.038	20.235	2.80	0	0	27.540	33.569				
		E	18.626	-13.222	18.423	2.30	0	0	27.565	33.550				
		F	26.623	6.326	19.998	2.65	0	0	27.080	32.873				

Table 2. Registration of Point Cloud to Point Cloud.

	Norm Used	Final Registration Values						ASD (cm)	RMS (cm)	Volume (m ³)	Volume Error (%)	Box Height (cm)	Box Ht. Error (%)	
		Δx (cm)	Δy (cm)	Δz (cm)	$\Delta \phi$ (°)	$\Delta \epsilon$ (°)	$\Delta \theta$ (°)							
Start at ($\Delta x, \Delta y, \Delta z$) = (0, 0, 0) (Cold Start)	ASD	C	0	0	0	0	0	0	NA	NA	1.861	9.53	93.3	2.03
		D	-7.889	17.475	-0.256	0.35	0	0	16.910	25.728				
		E	20.959	12.599	-1.456	0.95	0	0	14.020	25.250				
		F	12.406	18.133	0.504	0.60	0	0	20.044	32.902				
	RMS	C	0	0	0	0	0	0	NA	NA	2.304	35.59	92.6	1.27
		D	-8.122	17.584	-2.904	0.25	0	0	17.684	25.629				
		E	23.224	12.994	-7.367	1.30	0	0	16.829	25.761				
		F	14.741	31.937	-7.049	1.20	0	0	18.466	26.677				
Start off origin (warm start)	ASD	C	0	0	0	0	0	0	NA	NA	1.859	9.42	93.0	1.71
		D	-7.829	17.43	-0.267	0.35	0	0	16.908	25.727				
		E	17.587	17.063	-1.521	0.10	0	0	14.123	25.330				
		F	10.21	22.372	0.508	-0.10	0	0	20.887	33.774				
	RMS	C	0	0	0	0	0	0	NA	NA	1.913	12.57	93.6	2.36
		D	-8.129	17.6	-2.839	0.25	0	0	17.661	25.629				
		E	16.261	3.19	1.348	0.05	0	0	17.335	27.595				
		F	12.499	22.903	1.773	0.10	0	0	17.140	25.873				