

Registration of Range Images in the Presence of Occlusions and Missing Data

Gregory C. Sharp¹, Sang W. Lee², and David K. Wehe³

¹ Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI 48109

² Department of Media Technology
Sogang University, Seoul, Korea

³ Department of Nuclear Engineering and Radiological Sciences
University of Michigan, Ann Arbor, MI 48109

February 25, 2002

Abstract

We describe a method for the registration of range images in the presence of occlusions and missing data. Our approach uses the missing data as an additional clue for matching. Each point in both images, valid or missing, is classified and scored within a maximum likelihood framework. Occlusions, shadows, noise, partially overlapping views, and outliers are all considered in within this framework. We found that our registration method can be used to register real scenes with complex geometry, discontinuities, missing data, and overlap as small as ten percent.

1 Introduction

When building 3D models of real world objects from multiple range images, the range images must be properly positioned and oriented within a common coordinate frame before they can be merged into a coherent CAD description. This process, called registration, is often accomplished using a method such as the iterative closest point algorithm, or a related method [2, 5, 19, 12, 7, 4, 16, 13]. These algorithms operate on two range images, and search for the rigid transformation that best maps one view into another according to distance between the two surfaces.

The presence of occlusions and missing data confounds these algorithms because they use point-to-point or point-to-plane correspondences to align the data, and there is difficulty

in deciding which points can be matched safely, and which should be rejected. One solution to this problem is distance thresholding. According to the distance thresholding criteria, a match can be made if the distance from the point to the matched point is less than some threshold [19]. Another solution to this problem is view thresholding. According to the view thresholding criteria, a match can be made if a point in one view would be visible to the other sensor [3, 12]. A third possibility is to perform between pairs of correspondences [6]. Neither distance thresholding nor thresholding based on the consistency of differences are able to identify and exclude occluded regions from the matching process. Some view thresholding methods use geometric reasoning to identify occlusions, but they have difficulty with complex scenes that contain missing data.

The goal of this work is the design of a fully automatic registration method that will perform reliably on view pairs that are partially overlapping, contain occlusions, noise, outliers, and missing data. Further, a large amount of overlap should not be required, and an initial estimate should not be required. We have achieved all of these goals by using a maximum likelihood framework that reasons not only about the likelihood of receiving a given depth reading, but also reasoning about the likelihood of receiving no data. Missing data is expected to be correlated between views, as is the case when the target surface is non-reflective or very distant. Geometric reasoning is used to detect not only occlusions and out of view pixels, but also shadows cast by a structured light source. A robust estimator for valid point matches allows reliable registration in the presence of Gaussian and spike noise. A two stage search procedure is used to optimize the likelihood calculation. We demonstrate this registration method on on structured light range images of scenes with complex occlusions, discontinuities, and shadows.

This paper is organized as follows. In section 2, we describe a method that segments the range images into different classes, and assigns likelihood measures for each class. In section 3, we describe how the maximum likelihood search may be conducted, and in section 4 we demonstrate the effectiveness of this approach. Finally, in sections 5 and 6 we compare this approach with other approaches, and suggest future directions for this work.

2 Registration Quality Measure

2.1 Probabilistic Framework

The registration quality measure we use is a probabilistic measure which attempts to evaluate the likelihood of seeing a given scene. Suppose we have views \mathcal{S} and \mathcal{M} (the scene and the model), and the rigid transformation Θ brings view \mathcal{S} into alignment with view \mathcal{M} . The likelihood measure is the conditional probability

$$P(\mathcal{S}, \mathcal{M} | \Theta). \tag{1}$$

When $P(\Theta)$, $P(\mathcal{S}, \Theta)$, and $P(\mathcal{M}, \Theta)$ are each independent of Θ , the likelihood can be maximized by maximizing

$$P(\mathcal{S} | \Theta, \mathcal{M}) P(\mathcal{M} | \Theta, \mathcal{S}). \tag{2}$$

The quantity $P(\mathcal{S} | \Theta, \mathcal{M})$ is found by matching one view, the scene \mathcal{S} , against the other view, the model \mathcal{M} . We assume that each point \mathbf{s} in \mathcal{S} can be matched against the model

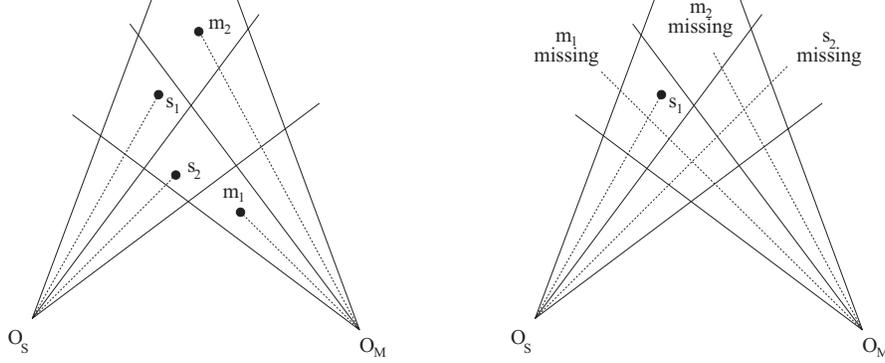


Figure 1: (Left) Example model and scene. (Right) Example model and scene that have missing points.

independently, so that we can minimize the negative log likelihood

$$-\log P(\mathcal{S}|\Theta, \mathcal{M})P(\mathcal{M}|\Theta, \mathcal{S}) = \sum_{\mathbf{s} \in \mathcal{S}} -\log P(\Theta(\mathbf{s})|\mathcal{M}) + \sum_{\mathbf{m} \in \mathcal{M}} -\log P(\Theta^{-1}(\mathbf{m})|\mathcal{S}). \quad (3)$$

Using the above assumptions, we proceed by quantifying the probability of viewing each point \mathbf{m} in \mathcal{M} and each point \mathbf{s} in \mathcal{S} for configuration Θ . To simplify the notation, we shall assume for the remainder of the paper that \mathcal{S} has been properly transformed by Θ , and when we refer to a scene point \mathbf{s} we are actually referring to a transformed scene point $\Theta(\mathbf{s})$.

2.2 Notation and Point Properties

Intuitively, it is likely that we may sense a point \mathbf{s} if the other view also sensed a range point nearby. On the other hand, if the other view has a missing data reading at the coordinate of this range point, this event is still somewhat likely when \mathbf{s} is very distant, or unreflective, or occluded. We will now make this intuitive this concept concrete, by performing a segmentation of the image into different classes, and assigning likelihood measures that depend upon the classification of each point.

Each pixel of the range image will be called *valid* if the sensor has returned a proper range value, or it will be called *missing* if the sensor could not return a range value. We assume the model of a laser range finder, where each pixel senses within a cone with origin O_s , and some known solid angle. The pixel then encodes the distance to a point on the closest visible surface patch within this cone. We will say that the point \mathbf{s} in \mathcal{S} *matches* \mathbf{m} in \mathcal{M} when \mathbf{s} lies within the cone sensed by the pixel \mathbf{m} . For each point \mathbf{s} in \mathcal{S} , and for each point \mathbf{m} in \mathcal{M} , we define the following attributes:

- s.corr** The point \mathbf{m} in \mathcal{M} which \mathbf{s} matches in the sense that \mathbf{s} lies within the cone of pixel \mathbf{m} .
- s.invcorr** Of all points in \mathcal{M} which match \mathbf{s} in the sense that they lie within the cone of \mathbf{s} , **s.invcorr** is the one that lies closest to O_s . This condition implies that **(s.invcorr).corr** = \mathbf{s} .

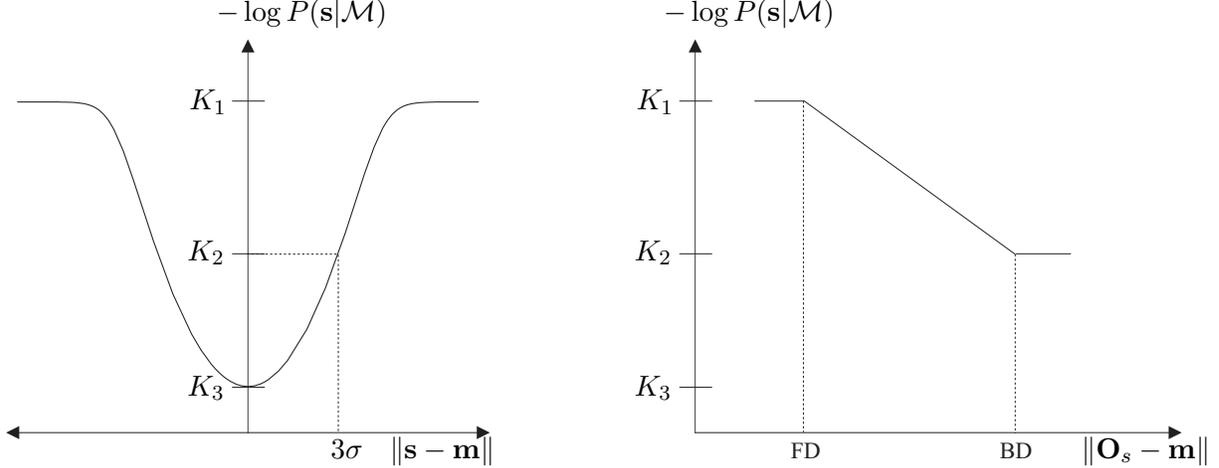


Figure 2: (Left) Cost function for fully matched points. (Right) Cost function for unmatched points.

If \mathbf{s} is missing, then $\mathbf{s}.corr$ is undefined. Further, if \mathbf{s} is valid, $\mathbf{s}.corr$ may be a missing data point in \mathcal{M} , or may be undefined if \mathbf{s} is out of the field of view of \mathcal{M} . To illustrate these definitions, consider figure 1 (left), where \mathbf{s}_1 , \mathbf{s}_2 , \mathbf{m}_1 , and \mathbf{m}_2 are all valid points. In this example, $\mathbf{m}_2.cor$ is \mathbf{s}_1 , $\mathbf{m}_1.cor$ is undefined, and $\mathbf{s}_1.cor$ and $\mathbf{s}_2.cor$ are both set to \mathbf{m}_1 . Further, $\mathbf{m}_1.invcorr$ is \mathbf{s}_2 , $\mathbf{s}_1.invcorr$ is \mathbf{m}_2 , and $\mathbf{m}_2.invcorr$ and $\mathbf{s}_2.invcorr$ are both undefined.

2.3 Fully Matched Points

In the case where both $\mathbf{s}.corr$ and $\mathbf{s}.invcorr$ are defined, we call this a *fully matched point*. A fully matched point corresponds to the event that our sensor has detected a surface patch that the other view has verified to exist. In figure 1 (left), \mathbf{s}_1 is fully matched because $\mathbf{s}_1.cor = \mathbf{m}_1$ and $\mathbf{s}_1.invcorr = \mathbf{m}_2$.

For fully matched points, we choose to model the probability $P(\mathbf{s}|\mathcal{M})$ using a mixed normal and uniform distribution, normal for correctly sensed points and uniform for outliers. With probability ρ_i the point is distributed normally with variance that depends upon the distance from the sensor, and with probability $1 - \rho_i$ the point is distributed uniformly over the range of valid outputs for the sensor. Let D_{\min} be the minimum sensing distance, or standoff distance, of the sensor. Let D_{\max} be the maximum sensing distance. Then, if $\mathbf{m} = \mathbf{s}.corr$, we have

$$-\log P(\mathbf{s}|\mathcal{M}) = -\log \left(\rho_i \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\|\mathbf{s}-\mathbf{m}\|^2}{2\sigma^2}} + (1 - \rho_i) \frac{1}{D_{\max} - D_{\min}} \right). \quad (4)$$

The shape of this cost function is shown in figure 2. It can be seen that the shape of the cost function is very similar to certain M estimators, and in fact, it is a variant of Cauchy's function [14]. The cost runs between $K_1 = -\log(\rho_i/(\sigma\sqrt{2\pi}) + (1 - \rho_i)/(D_{\max} - D_{\min}))$ and $K_3 = -\log((1 - \rho_i)/(D_{\max} - D_{\min}))$

Please note that we are taking some liberty with the probabilistic framework. Data points that are closer than D_{\min} or farther than D_{\max} are probability zero events, but we extend

the likelihood measure out to infinity because it leads to more predictable behavior when the values D_{\min} and D_{\max} are incorrect. In addition, we omit a constant term needed to normalize the joint distribution over different classifications. Since we are only computing relative likelihoods, this term is safely omitted.

2.4 Out of View Points

In the case where $\mathbf{s}.\text{corr}$ is undefined, we call this point *out of view*. This corresponds to the case where the point is out of the field of view of the camera. In figure 1 (left), \mathbf{m}_1 is out of view because $\mathbf{m}_1.\text{corr}$ is undefined.

When a point \mathbf{s} could not have been sensed by the other view \mathcal{M} , then \mathcal{M} contains no information about the possible values for \mathbf{s} , and $P(\mathbf{s}|\mathcal{M})$ is not a function of \mathcal{M} . Our method for dealing with this case is to assume the same probability as the fully matched at a distance $\|\mathbf{s} - \mathbf{m}\| = 3\sigma$. Hence, the cost for any out of view point is the constant labeled K_2 in figure 2:

$$K_2 = -\log P(\mathbf{s}|\mathcal{M}) = -\log \left(\rho_i \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{9}{2}} + (1 - \rho_i) \frac{1}{D_{\max} - D_{\min}} \right). \quad (5)$$

This choice for K_2 biases the registration to prefer matching points over leaving them unmatched. The 3σ error level is a convenient criteria, but K_2 may be increased or decreased to introduce more or less of this bias. In all cases K_2 should be between K_1 and K_3 .

2.5 Half Matched Points

In the case where $\mathbf{s}.\text{corr}$ is defined, and $\mathbf{s}.\text{corr}$ is valid, but $\mathbf{s}.\text{incorr}$ is undefined, we call this point *half matched*. One explanation for the half matched points is an occlusion in the scene that prevents the scene from sensing the model. Another explanation for the half matched point is a difference in sampling rates between the two images.

In figure 1 (left), \mathbf{m}_2 is half matched, because $\mathbf{m}_2.\text{corr} = \mathbf{s}_1$ and $\mathbf{m}_2.\text{incorr}$ is undefined, and \mathbf{s}_2 is half matched, because $\mathbf{s}_2.\text{corr} = \mathbf{m}_1$ and $\mathbf{s}_2.\text{incorr}$ is undefined. We may suppose that in the case of \mathbf{s}_2 the cause of the half matched point is that the surface between \mathbf{m}_1 and \mathbf{m}_2 was not sampled finely enough, so that no point could be matched with \mathbf{s}_2 . When a point is half matched, we assume that the point is occluded or undersampled, and assign the cost of K_2 described in equation 5, for the same reasons.

2.6 Unmatched Points

In the case where $\mathbf{s}.\text{corr}$ is defined, but $\mathbf{s}.\text{corr}$ is not a valid point in \mathcal{M} , and $\mathbf{s}.\text{incorr}$ is undefined, we call this point *unmatched*. The unmatched point corresponds to the event that we have sensed a point, but the other view provides no evidence that the point exists. In figure 1 (right), \mathbf{s}_1 is an unmatched point because $\mathbf{s}_1.\text{corr} = \mathbf{m}_1$ is missing, and $\mathbf{s}_1.\text{incorr}$ is undefined.

The likelihood of this event depends upon the sensor, but we model it as a low probability event, with exponentially increasing probability as the distance from the sensor increases. If

other information, such the surface normal direction or intensity is available, this information may also be included in the probability estimation. For the unmatched point, we have

$$-\log P(\mathbf{s}|\mathcal{M}) = \begin{cases} K_1 & \|\mathbf{O}_s - \mathbf{s}\| < D_{\min} \\ K_2 + \frac{\|\mathbf{s} - \mathbf{O}_s\| - D_{\min}}{D_{\max} - D_{\min}}(K_1 - K_2) & D_{\min} < \|\mathbf{O}_s - \mathbf{s}\| < D_{\max} \\ K_2 & \|\mathbf{O}_s - \mathbf{s}\| > D_{\max} \end{cases} \quad (6)$$

2.7 Fully Missing Points

In the case where \mathbf{s} is missing, and $\mathbf{s.invcorr}$ is undefined, we call this point *fully missing*. The fully missing point corresponds to the event that we could not sense the point from either view. In figure 1 (right), \mathbf{s}_2 and \mathbf{m}_2 are fully missing points.

This event often occurs due to the poor response of the sensor off of dark surfaces, or for surfaces which cannot be sensed because they are distant from the sensor for both viewpoints. Again, we have little information about this configuration, and we assign the cost of K_2 described in equation 5.

2.8 Half Missing Points

In the case where \mathbf{s} is missing, but $\mathbf{s.invcorr}$ is defined, we call this point *half missing*. This event corresponds to the case where we expected to see the surface described by $\mathbf{s.invcorr}$, but for some reason failed to receive a data reading. In figure 1 (right), \mathbf{m}_1 is a half missing points because \mathbf{m}_1 is missing, and $\mathbf{m}_1.invcorr = \mathbf{s}_1$.

The half missing point is a complementary case of an unmatched point, since $\mathbf{s.invcorr}$ must be an unmatched point in \mathcal{M} , and we treat it in a similar fashion. For $\mathbf{m} = \mathbf{s.invcorr}$, we assign

$$-\log P(\mathbf{s}|\mathcal{M}) = \begin{cases} K_1 & \|\mathbf{O}_S - \mathbf{m}\| < D_{\min} \\ K_2 + \frac{\|\mathbf{m} - \mathbf{O}_S\| - D_{\min}}{D_{\max} - D_{\min}}(K_1 - K_2) & D_{\min} < \|\mathbf{O}_S - \mathbf{m}\| < D_{\max} \\ K_2 & \|\mathbf{O}_S - \mathbf{m}\| > D_{\max} \end{cases} \quad (7)$$

2.9 Segmentation Example

To demonstrate the segmentation algorithm, we register view 1 against view 2 shown in figure 3. Figure 4 shows the resulting segmentation when the view pair is well registered, and figures 5 shows an example when the view pair is not well registered. When the alignment is poor, the number of unmatched and half missing points increases.

Notice the checkerboard effect in the segmentation of view 1. This is caused by a difference in the sampling rate between the two images. A method for correcting this effect is described in section 5. Also, notice how the occluded regions in view 2 are detected to the left of the front wall and to the left of the lamp in the center of the image.

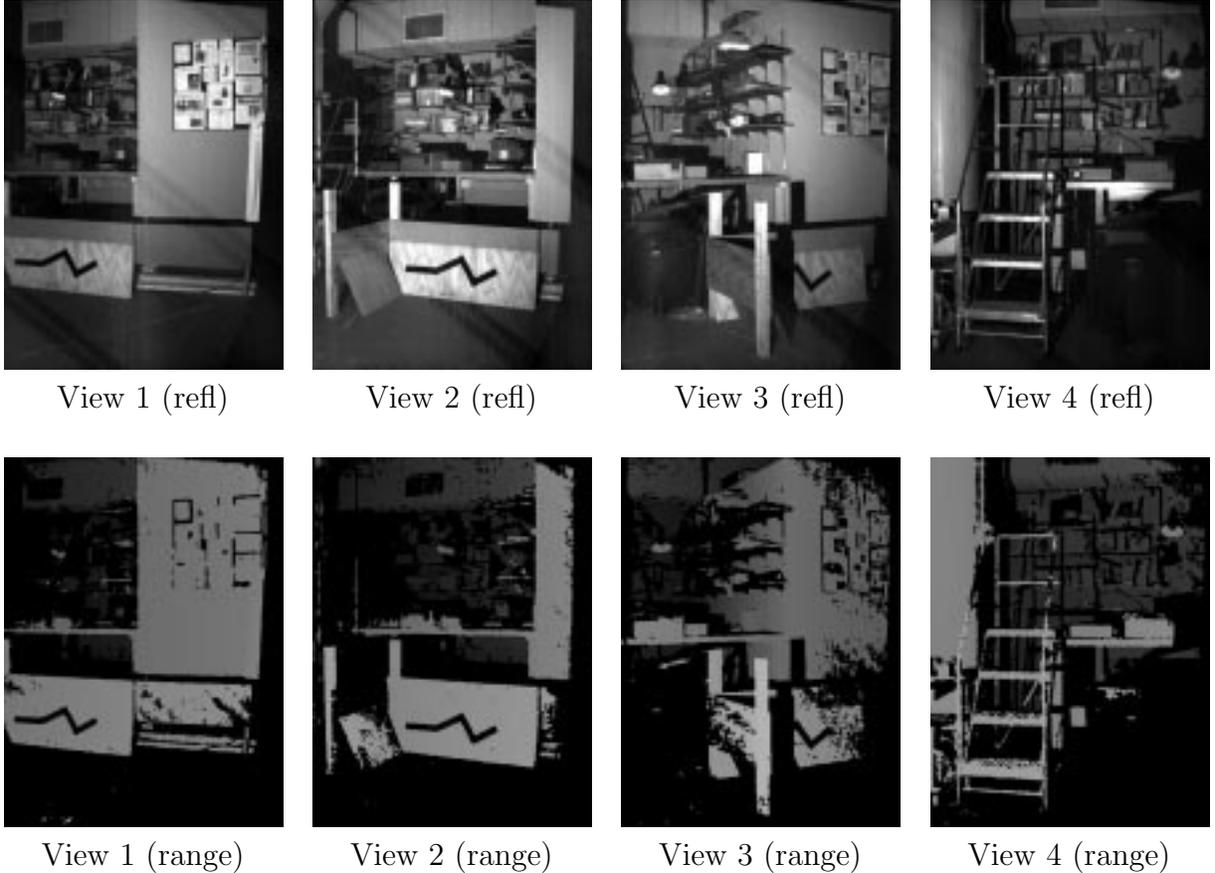


Figure 3: Range and reflection images of an indoor scene.

2.10 Multiple Centers of Projection

We can easily extend the concepts described in this section to include the case where our sensor has multiple centers of projection. For example, a structured light range sensor returns range values only for a surface patch that is visible to both the light source and the camera. Similar considerations arise for a stereo camera, which requires the surface patch to be seen from both cameras, but for this discussion we assume a structured light source. For these sensors, we extend the above segmentation procedure by computing the attributes $\mathbf{m.corr}$ and $\mathbf{m.invcorr}$ twice: once for the light source and once for the camera. Figure 4 shows the use of this technique for the detection of pixels that are out of the view of the projector used as a structured light source.

In addition, regions in the shadow of the light source, as shown in figure 6, can be identified. A valid pixel which is unmatched with respect to the camera, and either half matched or fully matched with respect to the light source is labeled *shadowed in the other image*. The converse case is when a pixel is half missing with respect to the camera, and either half matched or fully matched with respect to the light source. In this case, we label the pixel *shadowed in this image*. For both of these cases, we assign the probability $-\log P(\mathbf{s}|\mathcal{M}) = K_2$. An example of the detection of shadowed regions for view 2 and view 3 is shown in figure 7.



Figure 4: Segmentation results for view 1 and view 2 when they are properly registered. The color key for the bottom images is: (a) Fully matched: White, (b) Out of view (camera): Blue, (c) Out of view (projector): Cyan, (d) Half matched: Green, (e) Unmatched: Yellow, (f) Fully missing: Black, (g) Half missing: Red.



Figure 5: Segmentation results for view 1 and view 2 when they are not properly registered.

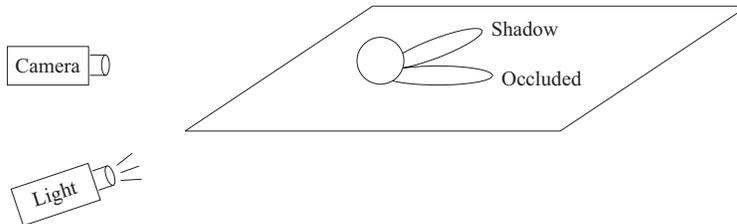


Figure 6: Shadows and occlusions in structured light range sensing.

3 Optimization

The registration procedure described above involves a difficult maximum likelihood search. The cost function is highly non-linear, it is discontinuous, and it has a multitude of tiny local minima. We have found a two stage optimization procedure to be useful for optimizing the registration criteria. The first stage performs a rough fit using a randomized search modeled on Masuda’s algorithm [12]. A small set of random samples are taken from the scene, and these samples are registered to the model using ICP. The solution of the ICP minimization is then tested against the maximum likelihood registration criteria, and if an improvement is shown, the new registration parameters are accepted. This idea is conceptually similar to RANSAC method [8, 4]. If an rough initial estimate of the pairwise registration is known,

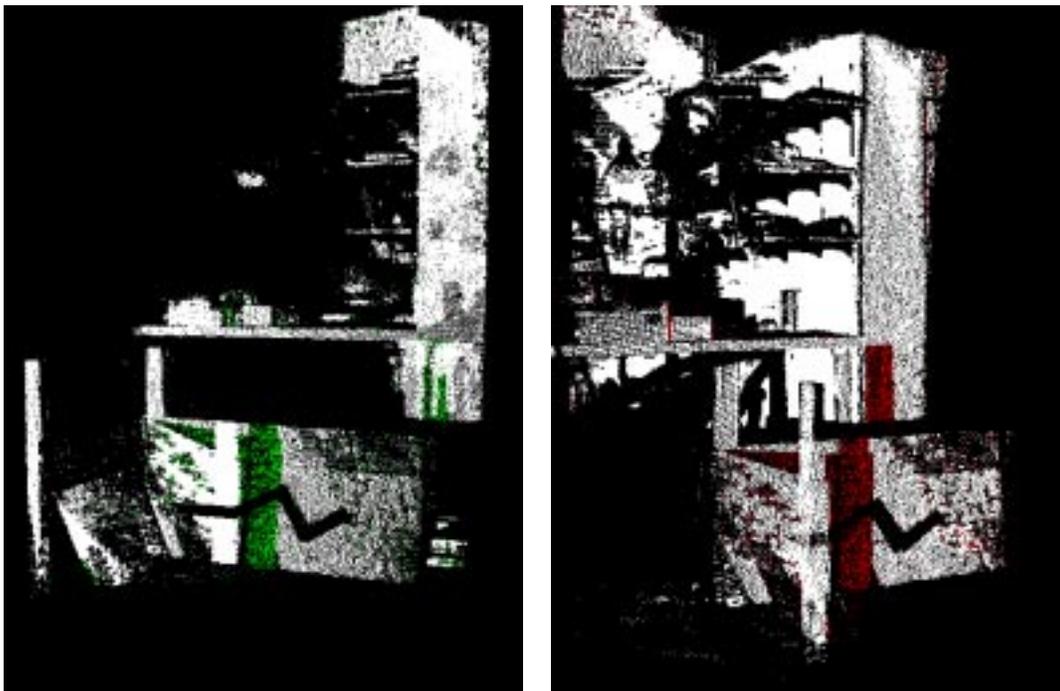


Figure 7: Segmentation results for view 2 (left) and view 3 (right). The data points that lie within the cast shadow of the structured light source in view 3 are identified in both images. The green pixels in view 2 are *shadowed in the other image*, and the red pixels in view 3 are *shadowed in this image*.

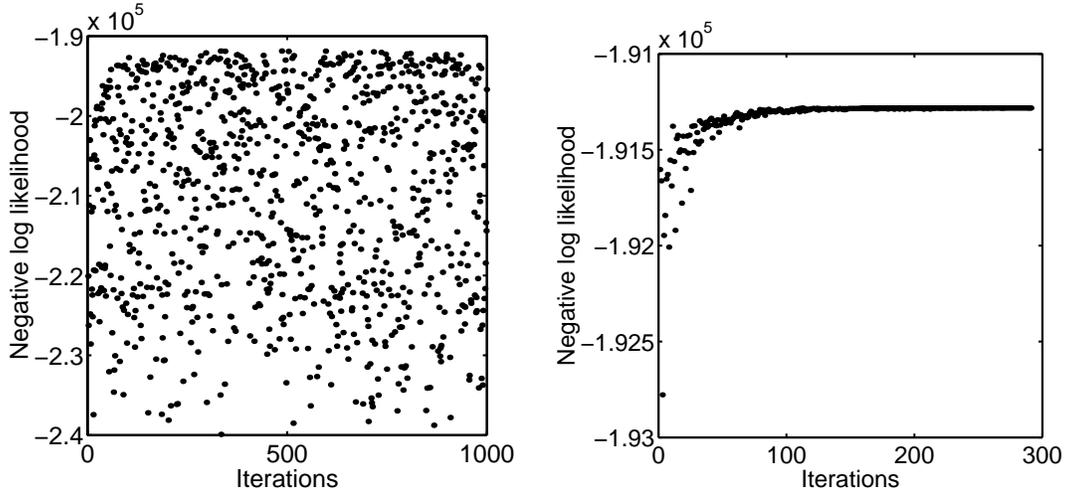


Figure 8: Negative log likelihood measurements for the first stage randomized search (left), and second stage Nelder-Mead search (right).

this first stage may be omitted.

The second stage uses a Nelder-Mead simplex search to refine the result [11]. Because the cost function has many discontinuities, the simplex search provides superior performance when compared to classical optimization methods such as conjugate gradient search. Once the initial estimate is known, convergence to a good local minimum is rapid.

It is also advisable to perform the optimization using multiple resolutions. First we run the randomized search on the subsampled image, increasing the size of cone used for matching by an appropriate amount so as not to introduce half-matched points. Then the Nelder-Mead search is run on subsampled images of increasing resolution. Since the purpose of the initial randomized stage is only to perform the rough alignment, there is no need to run that stage on higher resolutions.

4 Experimental Results

The range images of figure 3 were registered using the cost function described in section 2 and using the two stage search strategy described in section 3. The range images were taken from on board of a mobile robot in an indoor environment using a structured light range camera made from a DLP data projector and CCD camera. The method of Trobina was used to for calibration and signal processing of the sensor [17]. The parameter settings for minimum and maximum sensing distances, D_{\min} and D_{\max} , were set to 3 m and 8 m, the noise standard deviation parameter σ was set to 1 cm, and the inlier probability parameter ρ_i was set to 90%.

Convergence rates for the registration of view 1 and view 2 are shown in figure 8. For all scenes that we have tried, the first stage converges to an approximate solution in less than 1000 function evaluations, and the second stage converges in less than 400 function evaluations. On a 200 MHz Pentium II computer, one function evaluation on a 160×120

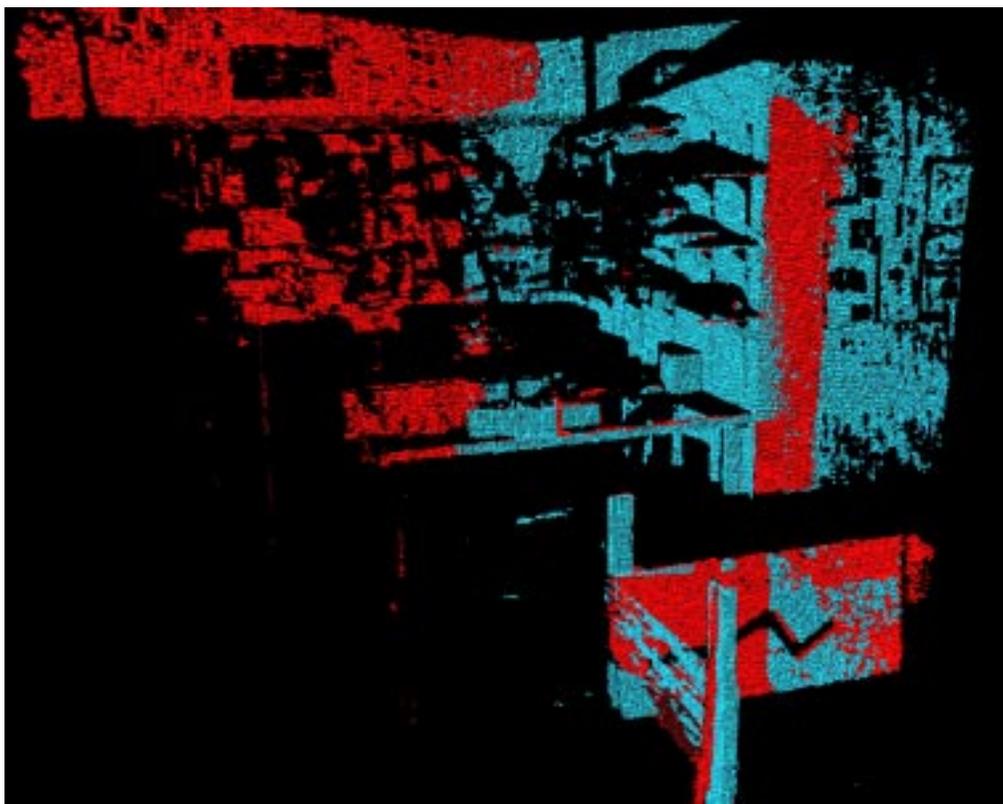
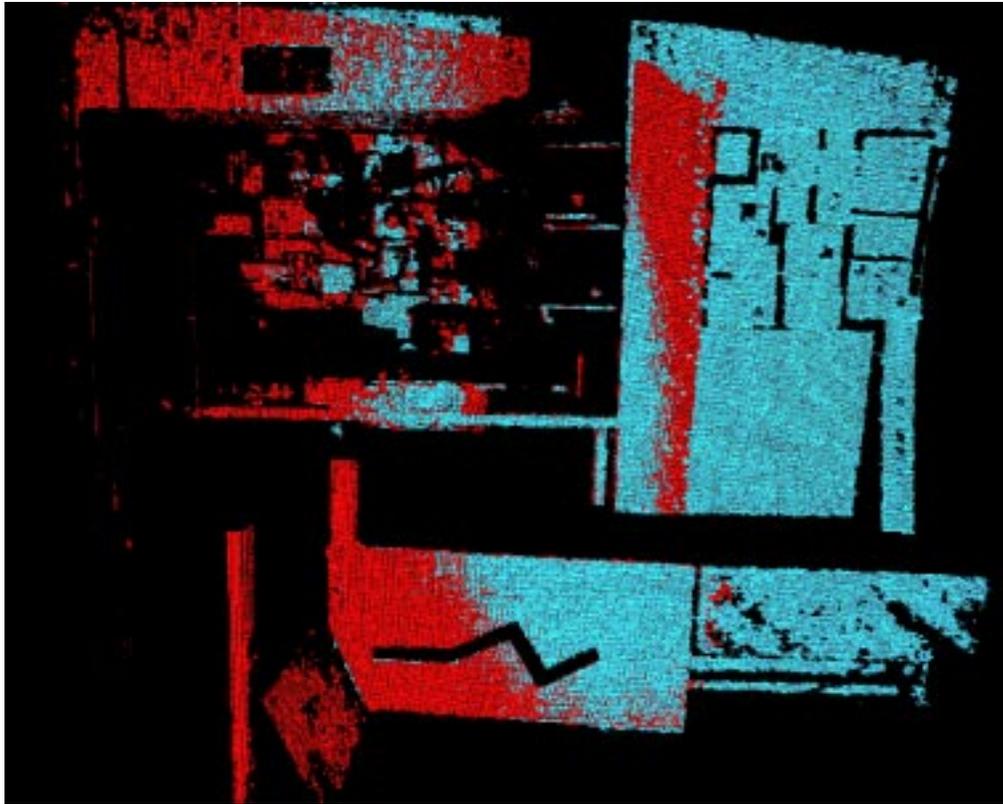


Figure 9: (Top) Registration of view 1 (blue) and view 2 (red) with 41.3% overlap. (Bottom) Registration of view 2 (red) and view 3 (blue) with 20.3% overlap.

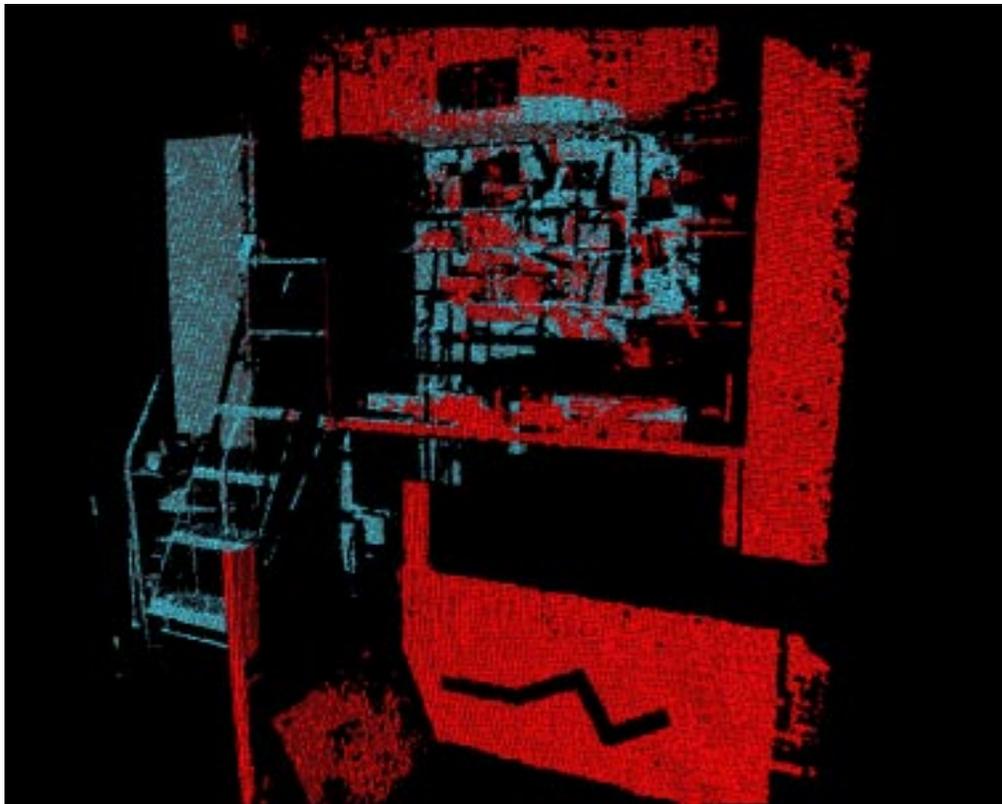


Figure 10: Registration of view 2 (red) and view 4 (blue) with 9.1% overlap.

image takes less than 300 ms, and a 640×480 image takes less than 500 ms.

Figures 9 and 10 shows the result of registering both view 1 against view 2, view 2 against view 3, and view 2 against view 4. The size of the overlapping region between the image pairs is approximately 41.3% for view 1 and view 2, approximately 20.3% for view 2 and view 3, and approximately 9.1% for view 2 and view 4. In each case, the correct registration was found fully automatically, without supplying an initial guess. To estimate our confidence about the quality of the fit, we compare the final log likelihood values for these three examples against the quantity K_2 . For this set of images, $K_2 = 5.4754$, while the average negative log likelihood measure is 4.9725 for the match between view 1 and view 2, 5.4039 for view 2 and view 3, and 5.4012 for view 2 and view 4. Since these measures are all less than K_2 , we feel some confidence that these matches are reasonable.

5 Discussion

Although many researchers favor using a threshold on the Euclidean distance between point matches to remove occlusions and outliers, we feel that geometric reasoning based on view thresholding is a better approach. It is, after all, a great performance gain to compute $\mathbf{s.corr}$ in constant time rather than searching through a tree in $O(\log n)$ average case time to find a closest point [3]. Furthermore, thresholding on Euclidean distance or distance consistency does not properly penalize registrations that are infeasible due to conflicts in visibility. For



Figure 11: Distance thresholding cannot recognize that this registration is inconsistent with sensor visibility.

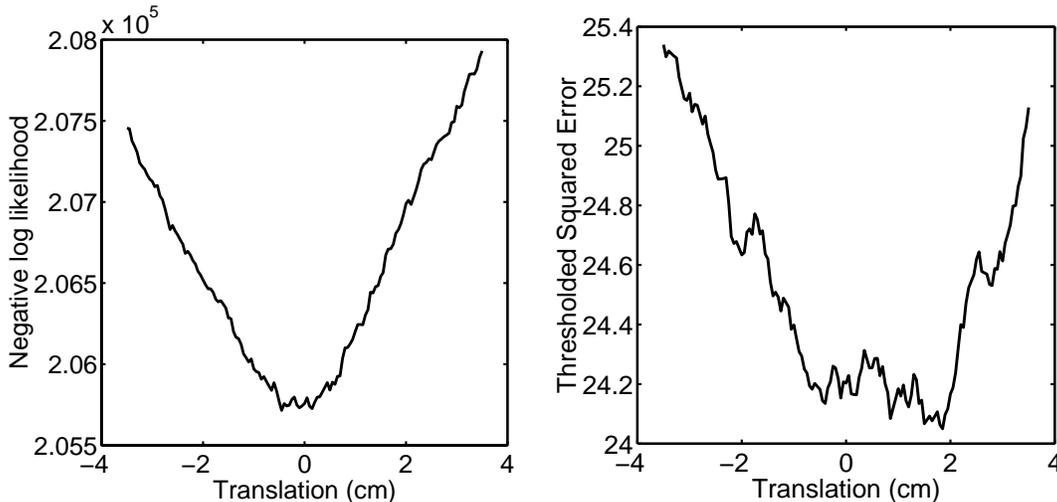


Figure 12: Registration quality measures for small translations near the best registration of view 2 and view 3. The view threshold method described in this paper (left), is less ambiguous criteria than the distance threshold criteria [19] (right).

example, figure 11 shows a registration that is very good using the distance thresholding, but is infeasible due to sensor geometry. Although matches at the boundary may be excluded by disallowing matches to boundary points [18], it is not obvious what constitutes a boundary for complex scenes with missing data. Even when view thresholding is performed, there has been much confusion about how to compare infeasible matches to those which are out of view. Almost always this is accomplished by requiring some minimum number of matches [3, 12], which implies that the amount of overlap should be known.

In addition, distance thresholding has been found in practice to have relatively poor localization. Figure 12 shows a comparison of the error surfaces for the maximum likelihood method described in this paper, and the normalized distance thresholding as described Zhang [19], for small translations near the correct registration. The method described in this paper is ambiguous to about 1 cm, while the distance thresholding method is ambiguous to about 3 cm.

It would be worthwhile to improve the segmentation results by using scan conversion rather than point projection to assign the value of $s.corr$. If s represents a surface patch, it rightfully may match several pixels in the model. This will greatly improve our ability to discriminate between the undersampling and occlusion for half matched points. Scan conversion will also greatly improve performance when a hardware Z-buffer is available [1].

As discussed earlier, the cost function for fully matched points is an M estimator, and

provides a method for limiting the bias of the estimate due to outliers. However, we find another place that the robust statistics can be used, which is in the estimate of $\mathbf{s.invcorr}$ itself. Currently, we are accepting that $\mathbf{s.invcorr}$ should be the closest point to the sensor, but when several points in \mathcal{M} match to \mathbf{s} , we will obtain a better estimate of the underlying surface using an M estimator.

Another important aspect of our algorithm is that it provides an absolute reference point to compare the quality of a match. As mentioned in section 4, we can compute the likelihood of two scenes that have no overlap, which assigns probability K_2 to all points. If a candidate match is less than, or not significantly better than this, we may wish to conclude that either the two scenes do not overlap or that the registration is low quality. This could be useful for building and verifying a graph of neighboring views used in multiview registration [15, 9].

Finally, we have not addressed the use of intensity information, or the use of surface normals. Intensity, and especially color, are crucial for the alignment of surfaces for which the shape has few remarkable features [10]. They will also be useful in the identification of missing data due the saturation of structured light sensors by specular reflection. The chief problem with using intensity information is the photometric calibration of the intensity response, which varies according to the distance to the surface, the surface normal, and the typically non-linear response of the sensor. Surface normals will be useful not only for computing the compatibility of the two surfaces, but also for computing the reliability or noise variance of the readings. Unfortunately, computing the surface normals is not a straightforward process since most range sensors exhibit considerable fine grain noise.

6 Conclusion

This paper has presented a fully automatic method for aligning range scans that contain occlusions and missing data. Both valid and missing range data are used to assess the quality of the match within a probabilistic framework. Using this framework, the registration of complex, noisy scenes with occlusions, missing data, and as little as 10% overlap have been demonstrated.

Acknowledgments

This material is based upon work supported by DOE under Award No. DE-FG04-86NE37969.

References

- [1] R. Benjemaa and F. Schmitt. Fast global registration of 3d sampled surfaces using a multi-z-buffer technique. *Image and Vision Computing*, 17(2):113–123, February 1999.
- [2] 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, February 1992.

- [3] G. Blais and M.D. Levine. Registering multiview range data to create 3d computer objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):820–824, August 1995.
- [4] C.S. Chen, Y.P. Hung, and J.B. Cheng. Ransac-based darces: A new approach to fast automatic registration of partially overlapping range images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1229–1234, November 1999.
- [5] Y. Chen and G.G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.
- [6] C. Dorai, G. Wang, A.K. Jain, and C. Mercer. Registration and integration of multiple object views for 3d model construction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(1):83–89, January 1998.
- [7] C. Dorai, J. Weng, and A.K. Jain. Optimal registration of object views using range data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(10):1131–1138, October 1997.
- [8] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, June 1981.
- [9] D.F. Huber. Automatic 3d modeling using range images obtained from unknown viewpoints. In *International Conference on 3D Digital Imaging and Modeling*, pages 153–160, May 2001.
- [10] A.E. Johnson and S.B. Kang. Registration and integration of textured 3d data. *Image and Vision Computing*, 17(2):135–147, 1999.
- [11] J.C. Lagarias, J.A. Reeds, M.H. Wright, and P.E. Wright. Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM Journal of Optimization*, 9(1):112–147, 1998.
- [12] T. Masuda and N. Yokoya. A robust method for registration and segmentation of multiple range images. *Computer Vision and Image Understanding*, 61(3):295–307, May 1995.
- [13] K. Pulli. Multiview registration for large data sets. In *International Conference on 3D Digital Imaging and Modeling*, pages 160–168, 1999.
- [14] W.J.J Rey. *Introduction to Robust and Quasi-Robust Statistical Methods*. Springer-Verlag, 1983.
- [15] G.C. Sharp, S.W. Lee, and D.K. Wehe. Toward multiview registration in frame space. In *IEEE International Conference on Robotics and Automation*, 2001.
- [16] G.C. Sharp, S.W. Lee, and D.K. Wehe. Invariant features and the registration of rigid bodies. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1):90–102, January 2002.

- [17] M. Trobina. Error model of a coded-light range sensor. Technical Report BIWI-TR-164, Communication Technology Laboratory, Image Science Group, ETH Zurich, 1995.
- [18] G. Turk and G. Levoy. Zippered polygon meshes from range images. In *SIGGRAPH*, pages 311–318, 1994.
- [19] Z.Y. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):119–152, October 1994.