Interactive 3D Scanning Without Tracking

Matthew J. Leotta, Austin Vandergon, Gabriel Taubin
Brown University
Division of Engineering
Providence, RI 02912, USA
{matthew_leotta, aev, taubin}@brown.edu

Abstract

Using inexpensive and readily available materials–a calibrated pair of cameras and a laser line projector–a 3D laser range scanner which requires no tracking is implemented in this paper. We introduce a planarity constraint for reconstruction, based on the fact that all points observed on a laser line in an image are on the same plane of laser light in 3D. This plane of laser light linearly parametrizes the homography between a pair of images of the same laser line, and this homography can be recovered from point correspondences derived from epipolar geometry. Points visible from multiple views can be reconstructed via triangulation and projected onto this plane, while points visible in only one view can be recovered via ray-plane intersection. The use of the planarity constraint during reconstruction increases the system's accuracy, and using the planes for reconstruction increases the number of points recovered. Additionally, an interactive scanning environment is constructed, where incremental reconstruction is used to provide continuous visual feedback. Reconstructions with this method are shown to have higher accuracy than standard triangulation.

1. Introduction

Laser range scanners (such as those offered by ShapeGrabber [1] and CyberWare [2]) provide an excellent way to recover shape data. Typically, however, a tradeoff exists between system cost and scan accuracy. High end commercial systems come with a substantial price tag, and require constant calibration and maintenance to recover consistent shape data. Some sort of tracking, such as the "articulated arm" in the Perceptron ScanWorks package [3] or FASTRACK in the Polhemus FastScan [4] system, is also usually necessary. Our system, by contrast, was conceived with low cost and simple scanning in mind. The setup includes a synchronized stereo pair and hand-held straight line projector, as well as a display where visual feedback in the form of incremental 3D reconstruction is provided. Figure 1 illustrates the configuration of the system. The cameras are fully calibrated with respect to the world coordinate system, and the fields of view of the two cameras overlap. The intersection of these two fields of view is the working volume where three-dimensional data can be captured. The line projector generates an arbitrary and unknown plane of light, which intersects the working volume and generates curves which should be visible in both images.

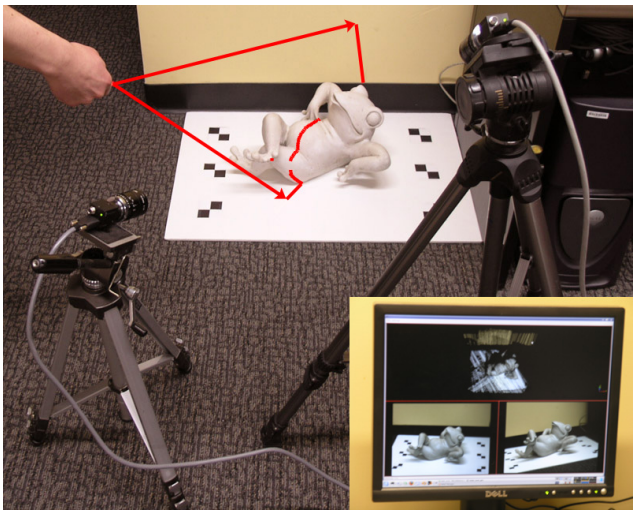

Figure 1. The scanning system, with a frog model in place.

The same physical setup is suggested by Davis and Chen [11] as a greatly improved laser range scanner. It maintains many of the desirable traits of other laser range scanners (accuracy, robustness) while eliminating actuated

components, thereby reducing calibration complexity and concomitantly increasing maintainability and scan repeatability. Furthermore, the elimination of actuated components drastically reduces the cost of the device. While standard triangulation used by Davis and Chen [11] can only recover points in the working volume, our approach can recover additional points visible from only one view. Our extended working volume is the union of the fields of view as illustrated in Figure 2.

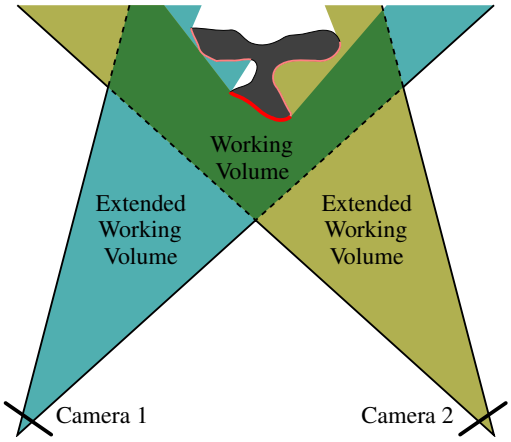

Figure 2. A top view of the working volume and scannable surfaces of a T-shaped object.

The rest of this paper is organized as follows: Section 2 discusses related work. Section 3 derives the equations for the novel reconstruction method. In Section 4, this method is applied in the implementation of an interactive scanning system. Section 5 provides error analysis and scanning results, and Section 6 concludes with a summary and future work.

2. Related Work

Slit scanners have been continuously improving since their development early in the history of range scanners. Section 2.2 of F. Blais' survey on the subject [7] provides a good overview of slit scanners' associated advantages and disadvantages. Work on this particular system started as an extension of a real time implementation of the shadow scanner created by Bouget and Perona [8]. This shadow scanner works using two calibrated planes on which the shadow cast from a stick illuminated by a point light source can be observed. By estimating the location of the front edge of the shadow plane in space, a reconstruction of an object can be obtained using a single camera via ray-plane intersection. Notably, in section 5, they suggest extending the shadow scanner to integrate multiple views; this suggestion is made, however, with the intention of integrating multiple scans easily to obtain a full 3D model of an object. Bouget and Perona also expanded upon their previous work in [10], using an arbitrary (rather than predictably oriented) set of shadow planes to eliminate the necessity of a background plane when doing reconstructions. Merging scans was addressed in [9], but the simultaneous use of multiple views and arbitrary shadows is not explored in their literature.

Subsequent research revealed that the same system setup we planned to test had been used by Davis and Chen [11]. In this paper, reconstructions were done by recovering point correspondences using epipolar lines, then doing triangulation. This paper provided both confirmation that a stereo pair with an uncalibrated light source was indeed capable of accurate reconstruction using triangulation alone, as well as a good reference concerning important implementation details such as laser line detection. Another system involving a plane of laser light was recently constructed by Zagorchev and Goshtasby [18], in which depth calculation was done via ray-plane intersection. The reconstructions were done from single views, however, and integrated using a reference frame which was also used to recover the laser plane. A third system, in principle basically the same as Bouget and Perona's shadow scanner with an added fast surface registration technique, was presented by Winkelbach *et al.* in [17].

Also closely related is the work done by Trucco and Fisher [16], which discusses many constraints on reconstructed points in a system with two cameras and a plane of laser light, including a depth constraint on corresponding points from two images; coplanarity is unmentioned in conjunction with this constraint. Their system relies on actuated components–specifically, a sliding table–and uses a "direct calibration" method to skirt around the use of camera models and plane parametrization. Direct calibration simply involves using an object of known geometry to calibrate the entire system, rather than calibrating the cameras and the plane individually, and therefore retains many of the disadvantages (such as constantly required recalibration) of complex laser range scanners.

A final group of papers [14, 6, 15, 5] describe methods for 3D reconstruction which either assume or take advantage of planarity in the observed scene. These are typically planar objects, however, such as tabletops or faces of buildings; none are planar illumination systems such as a laser line projector. In a paper by Szeliski *et al.* [15], a planarity constraint is used for reconstruction by limiting reconstructed points within detected planes in the image to be on those planes. Notably, insignificant additions to accuracy are reported due to their planarity constraint.

3. Reconstruction

Using the pinhole model for image formation, the equation of projection of a three-dimensional point p onto image point u in homogeneous coordinates is

$$\lambda u = K(Rp + T)$$

where λ is a non-zero scalar value, K is an upper triangular 3×3 matrix, R is a 3×3 rotation matrix, and T is three-dimensional translation vector. K, R, and T are all parameters of the camera. For the remainder of the paper it is assumed that all cameras are calibrated. That is, K, R, and T are known for each camera.

Since K is known, all points in pixel coordinates u can be converted to normalized image coordinates u' and

$$\lambda u' = \lambda K^{-1} u = Rp + T$$

To simplify notation, all image measurements will refer to the normalized coordinates. Hence, for a pair of cameras and a common point, the image formation equations become:

$$\begin{cases} \lambda_1 u_1 &= R_1 p + T_1 \\ \lambda_2 u_2 &= R_2 p + T_2 \end{cases}$$

For more information on image formation and camera calibration the reader is referred to [13].

The unknown plane of light is written as follows

$$\Pi = \{p : n_1 p_1 + n_2 p_2 + n_3 p_3 + n_4 = 0\}$$

where the coefficient vector $[n_1 \ n_2 \ n_3 \ n_4]^t$ is non-zero. Since this coefficient vector is determined up to a multiplicative scale factor, the family of three-dimensional planes has three degrees of freedom. An alternative vector notation is also used in this paper:

$$\Pi = \{p : n^t p - d = 0\}$$

where $n = [n_1 \ n_2 \ n_3]^t$ is a unit vector and $d = -n_4$ is the distance from the plane to the origin of the coordinate system.

3.1. Planar Curves and Homographies

If an object is placed inside the working volume, the set of points on the object illuminated by the line projector form a three-dimensional curve, C (See Figure 3). As a result of depth discontinuities the curve may be composed of various disconnected segments. However, the entire curve is planar and lies in the plane Π. As a result, the two image curves, c_1 and c_2, captured by the pair of cameras are related by a homography, H. This homography is the composition of the two perspective homographies; one from the first image plane to the plane of light, H_1^{-1}, followed by a second one from the plane of light to the second image plane, H_2. Since this homography is parametrized by the plane of light, the family of homographies produced by this process has only 3 degrees of freedom instead of the 8 degrees of freedom of a general unconstrained homography. Further information on plane induced homography constraints is available in Chapter 12.1 of [13].

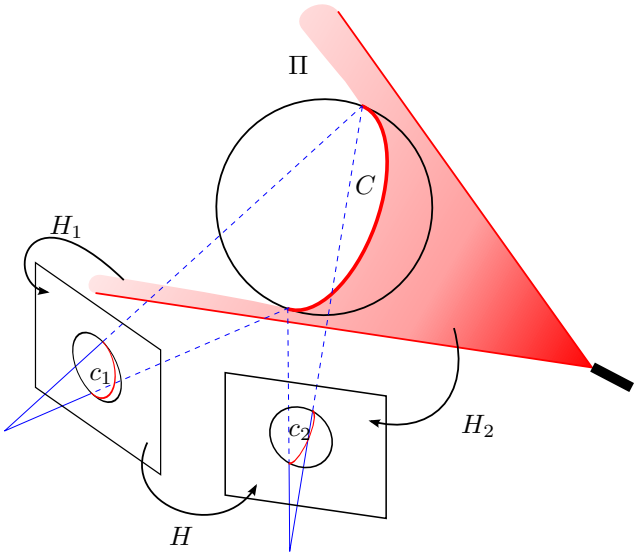

Figure 3. Homographies between the laser plane and image planes.

3.2. Planarity Constraint

Since the stereo pair is calibrated, corresponding points in the two image curves can be determined using epipolar geometry. By traditional triangulation each pair of corresponding image curve points determines a point on the three-dimensional curve. This is the point which minimizes the sum of the square distances to the two rays back projected through the image points. Due to measurement noise, these estimated points do not satisfy the co-planarity constraint. However, the constraint imposed by the reduced numbers of degrees of freedom in the parametrization of the homography allows for the estimation of the plane of light. Subsequently the location of points on the three-dimensional curve are estimated under the constraint of belonging to the estimated plane. This method produces more accurate results, as demonstrated in Section 5.

3.3. Homography Parametrization

The first step is to derive an expression for the homography $\xi u_2 = H\, u_1$ which transforms points on the first image plane onto points on the second image plane. Geometrically, this homography is defined by the following steps: 1) given a first image point u_1, compute the intersection point p of the ray corresponding to the image point u_1 with the plane Π; and 2) compute the image point u_2 as the projection of the point p onto the second image plane.

Algebraically, a point along the ray corresponding to u_1 can be written as

$$p = \lambda_1\, v_1 + q_1 \ , \tag{1}$$

in world coordinates, for a direction vector $v_1 = R_1^t u_1$, a center of projection $q_1 = -R_1^t T_1$, and some value of $\lambda_1 \neq 0$. For this point to be on the plane Π the following condition must be satisfied as well

$$0 = n^t p - d = \lambda_1\, (n^t v_1) + (n^t q_1 - d) \ , \tag{2}$$

or equivalently

$$\lambda_1 = \frac{d - n^t q_1}{n^t v_1} \ . \tag{3}$$

Note that the denominator is zero if and only if the ray defined by the first image point u_1 is parallel to the plane of light, which should never be the case here. Replacing the value of λ_1 just computed in the expression (1) for p we obtain

$$p = \frac{d - n^t q_1}{n^t v_1}\, v_1 + q_1 \ . \tag{4}$$

The projection of this point onto the second image is an image point u_2 which satisfies the projection equation

$$\lambda_2 v_2 = p - q_2 = \frac{d - n^t q_1}{n^t v_1}\, v_1 + q_1 - q_2 \ . \tag{5}$$

for some scalar value $\lambda_2 \neq 0$. Simple algebraic operations transform this expression into the following equivalent one

$$\xi\, v_2 = \left[(d - n^t q_1)\, I + (q_1 - q_2)\, n^t \right]\, v_1 \tag{6}$$

for some scalar value $\xi \neq 0$ ($\xi = \lambda_2\, (n^t v_1)$ in the previous expression), and I the 3×3 identity matrix. If A denotes the matrix inside the brackets, then the homography $\xi u_2 = H\, u_1$ which transforms points on the first image plane onto points on the second image plane is defined by the 3×3 matrix $H = R_2 A R_1^t$. Now, note that the matrix A can be written as a linear homogeneous combination of 3×3 matrices which only depends on the calibration parameters

$$A = n_1\, A_1 + n_2\, A_2 + n_3\, A_3 + n_4\, A_4 \tag{7}$$

with the coefficients of the plane of light as linear combination coefficients. As a result, so does H:

$$H = n_1\, H_1 + n_2\, H_2 + n_3\, H_3 + n_4\, H_4 \ . \tag{8}$$

Explicitly,

$$\begin{cases} A_1 &= (r_{11}^t T_1)\, I + (R_2^t T_2 - R_1^t T_1)\, e_1^t \\ A_2 &= (r_{12}^t T_1)\, I + (R_2^t T_2 - R_1^t T_1)\, e_2^t \\ A_3 &= (r_{13}^t T_1)\, I + (R_2^t T_2 - R_1^t T_1)\, e_3^t \\ A_4 &= -I \end{cases} \tag{9}$$

where r_{11}, r_{12}, r_{13} are the three columns of the rotation matrix $R_1 = [r_{11}\ r_{12}\ r_{13}]$ and e_1, e_2, e_3 are unit basis vectors (*e.g.* $e_1 = [1\ 0\ 0]^t$). Finally $H_j = R_2 A_j R_1^t$ for $j = 1, 2, 3, 4$.

3.4. Homography Estimation

Pairs of image points (u_1^j, u_2^j), $j = 1, \ldots, N$, corresponding to the same point on the three-dimensional curve are determined using epipolar geometry. For now, assume that each epipolar line intersects the imaged curve at exactly one point. Thus the corresponding image points are uniquely determined. The general case of epipolar matching is discussed in Section 4. Each of the image point pairs satisfy the homography equation $\xi_j u_2^j = H\, u_1^j$ for a different scale factor ξ_j. The scale factor is eliminated in the usual way (using a cross product), yielding two equations in H for each point pair:

$$\widehat{u_2^j}\, H\, u_1^j = 0 \tag{10}$$

where, if

$$u_2^j = \begin{bmatrix} u_{21}^j \\ u_{22}^j \\ 1 \end{bmatrix} \quad \text{then} \quad \widehat{u_2^j} = \begin{bmatrix} 1 & 0 & -u_{21}^j \\ 0 & 1 & -u_{22}^j \end{bmatrix} \ .$$

Equations (10) and (8) are combined to obtain the following matrix equation

$$\left[\widehat{u_2^j}\, H_1\, u_1^j \mid \widehat{u_2^j}\, H_2\, u_1^j \mid \widehat{u_2^j}\, H_3\, u_1^j \mid \widehat{u_2^j}\, H_4\, u_1^j \right]\, n = 0 \ . \tag{11}$$

Denote the 2×4 matrix within the brackets as L_j, and the $2N \times 4$ matrix resulting from vertically concatenating $L_1, \ldots, L_N$ as L. In the absence of measurement noise the linear equation $L\, n = 0$ should be satisfied, which implies that the matrix L should be rank-deficient, i.e. rank $(L) < 4$. The solution is unique if rank $(L) = 3$, which is the typical case. In practice there is measurement noise, and the solution is computed using the Singular Value Decomposition of the matrix L as the right singular vector corresponding to the minimum singular value. The second smallest singular value should be significantly larger than the minimum one.

If the points illuminated by the plane of light happen to be colinear, then a plane containing them is not uniquely determined. This singular case corresponds to rank $(L) = 2$. The handling of this degenerate case is discussed later on. Note, however, that the location of the three-dimensional points can still be estimated from triangulation.

3.5. Reconstruction of Point Locations

First consider the non-singular case when a unique plane of light Π can be determined. For each image point pair (u_1^j, u_2^j), a point p^j in three dimensions is determined as follows: 1) Compute a point $\hat{p}^j$ by standard triangulation. That is, compute the closest point to rays defined by u_1^j and u_2^j; and 2) Compute the point p^j as the orthogonal projection of $\hat{p}^j$ onto the plane Π. For points visible from only one camera (due to occlusion or indeterminate correspondence), use equations (1) and (3) to compute p^j as the intersection of the ray defined by u_1^j (or u_2^j) with the the plane Π.

Alternatively, points visible in both views could be reconstruct by ray-plane intersections independently. The resulting points could be averaged to produce a single reconstructed point. This method produces larger error than triangulation. The magnitude of the error in the reconstruction from each view depends on the angle between the ray and plane. As a result, one reconstruction may have much larger error than the other. Averaging these points with equal weight does not account for this and results in lower accuracy.

3.6. Singular Reconstruction

If the numerical rank of the matrix L is 2, the two right singular vectors associated with the two smallest singular values define two planes whose intersection is the line supporting all the illuminated points. Hence, the points lie on a line in three dimensions. Define the plane condition number κ_Π as the ratio of the second smallest to the largest singular value of L. When κ_Π is small, the rank approaches 2 and the plane estimate becomes unstable. Points visible in both views can still be reconstructed as before, even as the plane becomes degenerate. However, reconstructions from only one view are poor in this case. To avoid this, single view reconstructions are not computed when κ_Π exceeds a threshold. In practice, only points visible in both views are reconstructed when $\kappa_\Pi < 0.03$.

4. Scanner Implementation

Using the reconstruction technique in Section 3 we implemented an interactive scanning system described below. The system currently operates at about 4.5 frames per second and provides immediate iteratively updated reconstruction results. This allows the user to direct the laser to regions of sparse or missing data.

4.1. System Hardware

As in [11], a laser line projector is constructed from a cylindrical lens and a laser pointer. A small length of 6mm diameter Pyrex glass rod is attached to the front of an inexpensive consumer grade laser pointer. The quality of laser line generated is sufficient for our purposes. We found that a clear acrylic rod of the same dimensions resulted in a speckled line and was not suitable. However, using a stock glass rod for the lens was acceptable and a precision cylindrical or line generating lens was not required.

To capture a stereo sequence of images we used a pair of synchronized 1394b cameras capturing at a resolution of 1024×768. Each camera is mounted on separate tripod (Figure 1) so that the working volume can be adjusted for objects of various sizes. The cameras are positioned such that many points on the surface of the object are visible in both views while some additional points are visible only in one view. As with standard triangulation, a larger baseline results in more accurate reconstructions. The cameras may be calibrated in advance of scanning using any of the standard calibration techniques [13].

Our method is also compatible with the camera configuration used by Davis and Chen [11]. They use a single camera and mirrors capture two views of an object in a single image. This offers an even less expensive alternative to a synchronized stereo pair of cameras.

4.2. Line Detection and Matching

The laser line is detected independently in each image. A mean image is computed and subtracted from the current image. This removes the appearance from ambient lighting and leaves only the laser line. Gaussian smoothing and peak detection are applied to estimate the curve points. The points are then refined to subpixel location using parabolic interpolation. Finally the curve points are linked using standard edge linking to produce piecewise linear curves.

Correspondences are computed by rectifying the detected curves using homographies that map the epipoles to infinity, align the epipolar lines, and roughly preserve scale (refer to Chapter 10.12 of [13]). The curves are resampled at equal intervals in the rectified space resulting in a set of correspondences along epipolar lines. In practice, some points may have multiple correspondences. Davis and Chen [11] address this problem by discarding all ambiguous correspondences. The approach in this paper allows these ambiguities to be resolved by using the estimated homogra-

phy, H, resulting in a more detailed reconstruction. However, the homography parameters must first be estimated as in Section 3 using unique correspondences. The ambiguous correspondences are discarded for the purpose of homography estimation.

Estimating the homography from the remaining unique correspondences is still problematic. An ambiguous correspondence combined with an occlusion can produce a unique, but incorrect, correspondence. These incorrect correspondences are outliers and make up a small portion of the data in practice. Yet a single outlier can introduce a significant amount of error in a linear estimate. To compensate for this, RANSAC [12] is used. RANSAC is applied to the set of unique correspondences to get a robust estimate of the homography and discard outliers. RANSAC uses subsets of three correspondences to find many non-degenerate solutions to the estimation problem resulting from (11). Given a randomly selected subset of three correspondences the parameters of plane Π are computed and H is recovered as in (8). For each correspondence (u_1^j, u_2^j), the symmetric transfer error is computed:

$$E^j = \sqrt{(H\,u_1^j - u_2^j)^2 + (H^{-1}\,u_2^j - u_1^j)^2} \qquad (12)$$

where distances in (12) are computed using inhomogeneous coordinates. E^j is small for correspondences agreeing with H (inliers) and large for others (outliers). Applying a threshold to E^j classifies correspondence j as an inlier or outlier. In practice RANSAC is not very sensitive to the value of this threshold. The experiments in Section 5 uses a threshold value of 2. After RANSAC finds the solution with the most inliers, The same threshold procedure is applied to the ambiguous correspondences to find additional inliers. A final estimate of Π is computed using all of the inliers.

5. Evaluation and Results

To evaluate the accuracy of the algorithm, two simple objects of known dimension, a cylinder and a sphere, were scanned. The diameter of the cylinder is 3.125 inches (79.375 mm) and the diameter of the sphere is 4 inches (101.6 mm). Points visible in both views were reconstructed by triangulation and triangulation with the planarity constraint for comparison. Figure 4 shows the reconstructions by triangulation. The reconstructions using the planarity constraint have a similar appearance.

Cylinder and sphere models were fitted to the pairs of reconstructions using Gauss-Newton iteration to find a least-squares solution. The triangulated cylinder data had a diameter of 79.60 mm (0.28% off) and standard deviation from the estimated surface of 0.6299 mm. The plane restricted cylinder data set had a diameter of 79.13 mm (0.31% off) and a standard deviation of 0.5447 mm. While the error

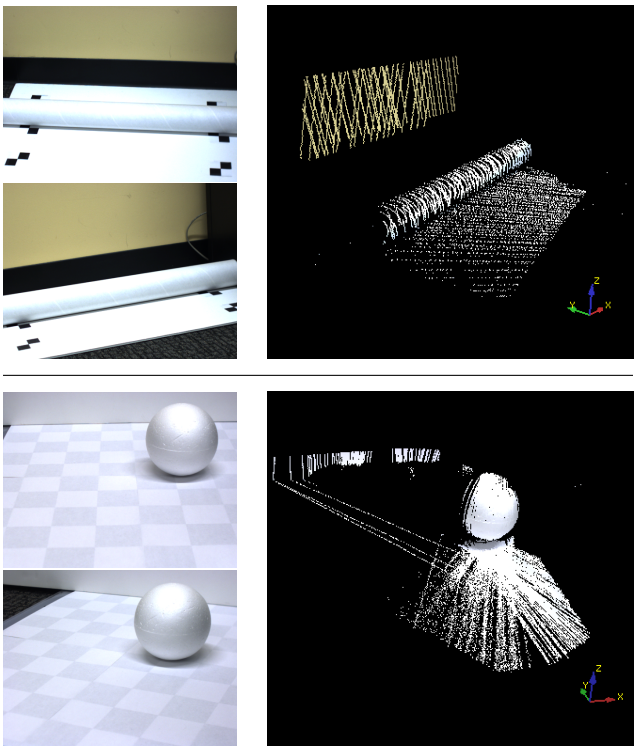

(a) Camera views (b) Triangulated points

Figure 4. Reconstructions of known objects used for evaluation

in diameter is essentially comparable, the standard deviation of the results obtained using the planarity constraint is 13.5% smaller. Figure 5 shows histograms of the two datasets.

The triangulated sphere data had a diameter of 101.77 mm (0.17% off) and a standard deviation of 1.825 mm. The plane restricted sphere data set had a diameter of 101.74 mm (0.14% off) and a standard deviation of 1.571 mm. Again the estimates are very accurate, and the application of the planar constraint reduces the standard deviation of the points by about 13.9%.

Figure 6 shows example reconstructions of various objects. Compared are triangulated points, points reconstructed from both views as described above, and these points combined with points reconstructed from a single view. The 3D points have been colored by projecting back into the mean ambient light image. Note the reduction of outliers from (b) to (c) and the increased number of points from (b) to (d). The top example (the frog model) results in 50,986 points from triangulation. Our method reconstructs 47,271 points from both views since many triangulated points are deemed to be outliers. An additional 46,734 and 10,483 points were reconstructed from only the

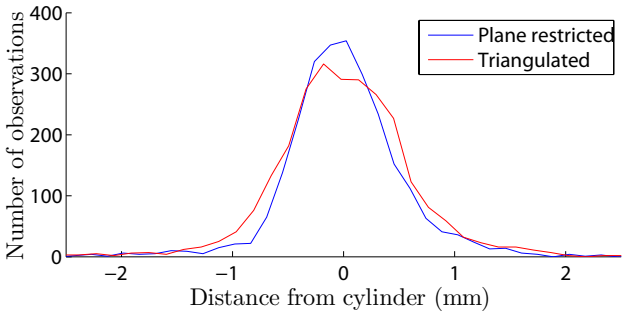

Figure 5. Histograms showing the distribution of distances from the corresponding fitted cylinder models.

right and left views respectively. In total, our method results in 104,488 points, more than twice that of standard triangulation.

6. Conclusions and Future Work

This paper has introduced a planar constraint for the 3D reconstruction of points in an inexpensive laser scanning system. The plane is estimated directly from the image data with no need for tracking the laser projector. The estimated plane is used to detect and remove outlier points, constrain the points to lie within the plane, and reconstruct additional points observed only by one camera. Applying the planar constraint to the triangulated points has been shown to reduce the standard deviation of the reconstruction error. After closer examination, it appears that the improved accuracy may be mostly due the elimination of outliers in the triangulated data. A preliminary investigation has indicated that in portions of the data the reconstruction accuracy is not significantly improved by projecting onto the plane. We will continue to investigate this issue and also evaluate other methods for 3D reconstruction using the planar constraint.

A natural extension of the work done here would be to take advantage of the fact that all reconstruction methods are viable for any number of cameras greater than or equal to two. One can envision a system containing multiple pairs of cameras arrayed around an object, where the single viewpoint reconstruction capability of this system was fully exploited in addition to triangulation with the planarity constraint to obtain very complete and accurate scans. Additionally, we may investigate automatic recovery of the extrinsic camera calibration parameters using the detected laser curves. This could be done using detecting features of the curves or simply by removing the cylindrical lens and detecting the laser point projected on the surface. This would simplify the calibration step currently required.

References

[1] http://www.shapegrabber.com/.

[2] http://www.cyberware.com/.

[3] http://www.perceptron.com/system_consist.html.

[4] http://www.polhemus.com/?page=Scanning_Fastscan.

[5] C. Baillard and A. Zisserman. Automatic reconstruction of piecewise planar models from multiple views. *CVPR*, 02:2559, 1999.

[6] S. Baker, R. Szeliski, and P. Anandan. A layered approach to stereo reconstruction. *CVPR*, 00:434, 1998.

[7] F. Blais. Review of 20 years of range sensor development. In *Videometrics VII. Edited by El-Hakim, Sabry F.; Gruen, Armin; Walton, James S. Proceedings of the SPIE, Volume 5013, pp. 62-76*, 2003.

[8] J.-Y. Bouguet and P. Perona. 3d photography on your desk. In *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, page 43, Washington, DC, USA, 1998. IEEE Computer Society.

[9] J.-Y. Bouguet and P. Perona. 3d photography using shadows in dual-space geometry. *Int. J. Comput. Vision*, 35(2):129–149, 1999.

[10] J.-Y. Bouguet, M. Weber, and P. Perona. What do planar shadows tell about scene geometry? *CVPR*, 01:1514, 1999.

[11] J. Davis and X. Chen. A laser range scanner designed for minimum calibration complexity. *3DIM*, 00:91, 2001.

[12] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.

[13] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, first edition, 2000.

[14] M. Irani, P. Anandan, and D. Weinshall. From reference frames to reference planes: Multi-view parallax geometry and applications. *Lecture Notes in Computer Science*, 1407:829–845, 1998.

[15] R. Szeliski and P. H. S. Torr. Geometrically constrained structure from motion: Points on planes. In *SMILE'98: Proceedings of the European Workshop on 3D Structure from Multiple Images of Large-Scale Environments*, pages 171–186, London, UK, 1998. Springer-Verlag.

[16] E. Trucco and R. B. Fisher. Acquisition of consistent range data using local calibration. pages 3410–3415, Los Alamitos, CA, USA, 1994. IEEE Computer Society.

[17] S. Winkelbach, S. Molkenstruck, and F. Wahl. Low-cost laser range scanner and fast surface registration approach. In *DAGM06*, pages 718–728, 2006.

[18] L. Zagorchev and A. Goshtasby. A paintbrush laser range scanner. *Comput. Vis. Image Underst.*, 101(2):65–86, 2006.

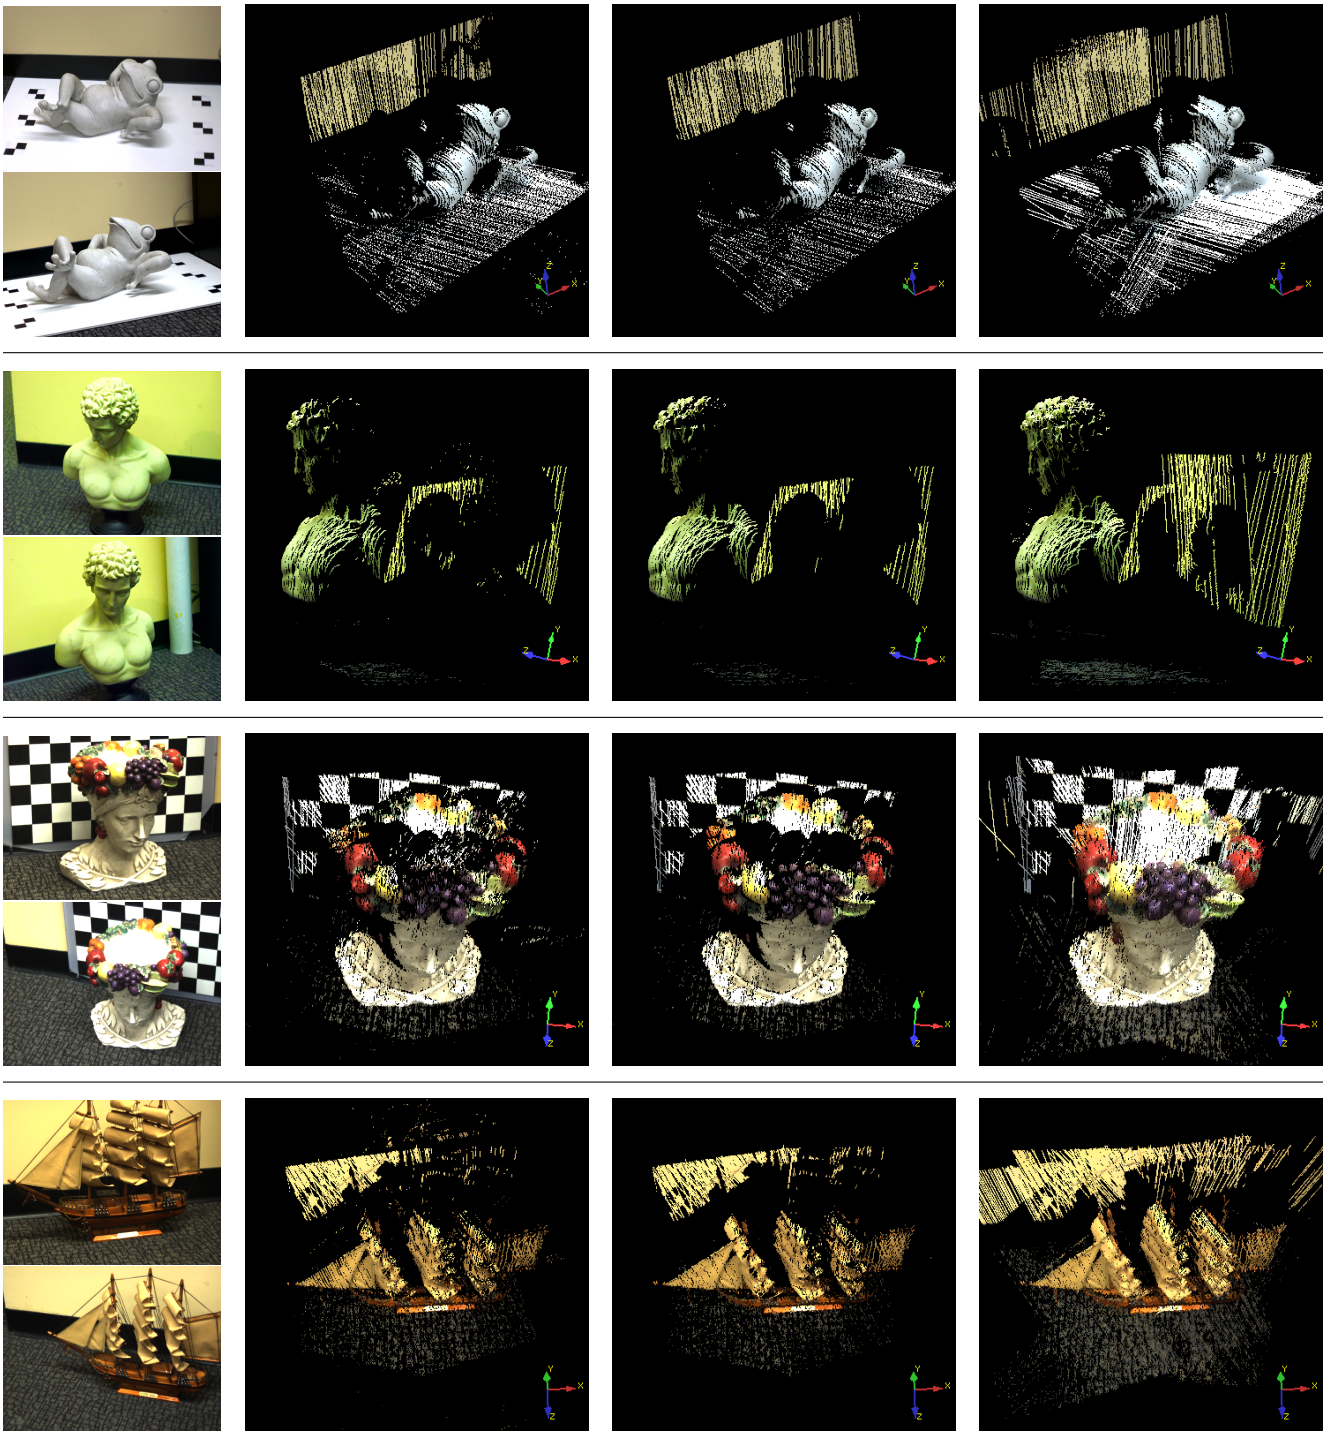

(a) Camera views (b) Triangulated points (c) Points from both views (d) All reconstructed points

Figure 6. Camera views and reconstruction results using triangulation and our method. Points in (c) are reconstructed from both views and (d) shows the addition of points seen only from one view.