

UC Davis

IDAV Publications

Title

Rotating Scans for Systematic Error Removal

Permalink

<https://escholarship.org/uc/item/8gx0d9gw>

Authors

Abbasinejad, Fatemeh

Kil, Yong J

Sharf, Andri

et al.

Publication Date

2009

Peer reviewed

Rotating Scans for Systematic Error Removal

Fatemeh Abbasinejad, Yong Joo Kil, Andrei Sharf, Nina Amenta

University of California, Davis

Abstract

Optical triangulation laser scanners produce errors at surface discontinuities and sharp features. These systematic errors are anisotropic. We examine the causes of these errors theoretically, and we study the correlation of systematic error with edge size and orientation experimentally. We then present a novel processing method for removing systematic errors, by combining scans taken at several different orientations. We apply an anisotropic filter to the separate scans, and use it to weight the data in a final combination step. Unlike previous approaches, our method does not require access to the scanner's internal data or firmware. We demonstrate the technique on data from laser range scanners by two different manufacturers.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.3]: Range Scanning—

1. Introduction

Progress in range scanning techniques has been driving a large amount of work in geometric modeling over the past two decades. There are many different methodologies and applications, including satellite and airplane stereographic scanning of the earth, LIDAR systems acquiring street views at driving speed and structured light scanners acquiring dynamic scenes. Nevertheless optical triangulation laser range scanning remains the most common technique for capturing surface data, at least in computer graphics.

The accuracy of optical triangulation hinges on the ability to locate the intensity peak of a beam of laser light reflected by the surface. Variations in surface reflectance and shape result in systematic errors in the captured depth map [SLPA90, BLS92, CL95]. The systematic error is anisotropic, as seen in Figure 1. The error is most pronounced at sharp surface edges that are perpendicular to the triangulation baseline (the direction separating the optical sensor and the laser emitter), and does not appear at surface edges running in the direction parallel to the baseline.

Curless and Levoy treated this problem, among others, in a pioneering paper [CL95] about improving the internal peak detection and triangulation firmware in a Cyberware scanner, using a “space-time” analysis which considered multiple charged coupled device (CCD) sensor frames when detecting the intensity peaks in each individual frame. Both of the modern scanners we experimented with, however, pro-



Figure 1: Removing systematic error from a scanned relief of a wall in Persepolis. In the close-up, upper right, we see an extra ridge to the right of the spear. Errors at depth discontinuity edges perpendicular to the triangulation baseline are typical of optical triangulation laser scanners. Below, capturing four depth maps at different orientations and combining them using our novel anisotropic filter removes these systematic errors.

duce anisotropic systematic errors which might have been eliminated by better internal processing. Since out-of-the-box scanners do display systematic error, we present a post-




Figure 2: *Left, a close-up of a depth map captured from our customized bulls-eye calibration object (see Figure 3 and Section 5), produced with a high-precision laser range scanner (Minolta Vivid 910). The corners should all be 90° . A side view of a strip cut along the scanner's sensor-emitter axis shows the systematic error: a stair-step error on each rising edge, and an overshoot before each falling edge. Right, combining four such depth maps successfully removes the systematic error.*

processing option which can be employed by the average user in search of better quality data.

One reason, perhaps, that removing systematic error is not a priority for scanner manufacturers is that it is not usually obvious in captured scans. This is because the systematic error is usually masked by larger, more random errors. Other sources of error include the difficulty of peak detection in the rasterized CCD image even in the absence of sharp edges, interpolation and quantization error, random noise in the CCD sensor and laser speckle, a physical phenomenon produced by interference between the reflected waves of the coherent light of the laser. Random error is comparatively easy to remove, as shown in recent work by Kil et al. [KMA06] (by averaging large amounts of data) or Diebel et al. [DTB06] (using a Bayesian approach). When random error is removed, however, the systematic error becomes far more obvious (see Figure 1 and Figure 2). In these situations the removal of systematic error becomes important.

Systematic errors also appear at sharp reflectance discontinuities as well as at the sharp depth discontinuities that we study here. These errors tend to be smaller, and are often avoided in practice by covering the object with a matte powder or paint, or scanning a cast rather than the original object. Depth discontinuities, of course, present a more fundamental problem that cannot be handled with physical work-arounds.

Our basic approach is to collect multiple scans using an out-of-the box scanner, differing by a rotation of the object around the scanner's z -axis. Each edge appears without systematic error in at least one of these orientations. We com-

bine the scans in software, weighting parts of the scan which we believe to be accurate heavily and giving low weights to the parts which we expect have systematic errors.

We apply this approach in two situations in which random error has already been removed from the input scans by averaging. The first context is the super-resolution technique of Kil et al. [KMA06]. Super-resolution is an image processing technique for combining many low-resolution images to make a single high-resolution image. They generalized super-resolution to laser scanner data; systematic error becomes more obvious as noise is removed and resolution is improved. In this paper, we very successfully remove systematic error from super-resolved depth maps made from data collected using a Minolta Vivid 910 scanner; Figure 1 and Figure 2 are examples of this application.

Additionally, we remove noise from scans produced by a "commodity" scanner made by NextEngine. The NextEngine is much less expensive than the Minolta Vivid 910 and the scans it produces are noisier. Hence, improving the depth maps produced by the NextEngine by post-processing in software would give a low-cost, high-quality scanning solution. Averaging only few scans does greatly reduce the noise; but again, when random noise is removed the systematic error becomes apparent. We apply our method to combine several de-noised depth maps, producing output that is dramatically better than the raw NextEngine scans, as seen in Figures 9 and 11.

2. Related work

The problem of accurately selecting the peak from the noisy, rasterized image of a reflected beam of laser light in the CCD sensor has received a fair amount of attention; for a survey see [FN96]. Several works [BR86, FSCP04] handle the noise and rasterization error by applying finite impulse response (FIR) filters to smooth the signal and increase the precision of the peak position; some averaging, smoothing or fitting is necessary because of the noise.

[BLS92] presents a thorough study of optical triangulation performance and measurement limitations. The authors note that at sharp edges, the detected peak shifts and produces an erroneous measurement of surface height. In [KGC91], the authors design a VLSI sensor array that gathers depth data in parallel as a scene is swept by a moving stripe. They notice that the parallel lightstripe method increases the accuracy of the depth measurement. In this setting the peak is detected using the additional time domain of the same position. Similarly [CL95] perform a specific space-time analysis by recording a sequence of corresponding CCD images across time. The improved signal is computed by using a Gaussian fit.

All of these approaches directly access the signals recorded on CCD sensors, or even explicitly alter the CCD device. In contrast, in this paper we treat the range scanner

as a black box, and process the point clouds that it outputs, a more realistic and useful approach for end users.

Diebel, Thrun and Bruenig [DTB06] proposed a computational method for noise reduction, based on a Bayesian formulation combining a model of the process introducing the error with a prior on the true shape of the surface. They assumed Gaussian error, which indeed removes random noise, but does nothing about systematic error. One could imagine developing a similar Bayesian algorithm using a computational model for systematic error, and indeed we considered this approach. But developing an accurate error model seems far from trivial, as we discuss in Section 3, and an inaccurate error model might introduce artifacts of its own (whereas the worst one fears from their Gaussian is over-smoothing).

Kil et al. [KMA06], in their super-resolution work, made the observation that systematic error becomes grossly evident when random sources of error are removed, and we use their technique to generate super-resolved scans.

More generally, combining data is an important idea in geometry acquisition, for example in the work of Nehab et al. [NRDR05], in which captured normals and surface position data were combined to greatly improve the resolution of the final models. Our use of scans taken from multiple orientations is also somewhat reminiscent of the fascinating work of Raskar et al. [RTF*04], in which flashes from multiple directions were used to extract edges in photographic images for non-photorealistic rendering.

In the following sections, we first describe our experimental observations of systematic error and discuss why this specific pattern of error occurs at sharp depth discontinuities. Then in Section 4, we present the algorithm for weighting and combining the scans. In Section 5 we discuss our data acquisition and pre-processing, and we present our results and experiments in Section 6, and we conclude by describing future directions and limitations.

3. Systematic Error

In an optical triangulation scanner, the laser emitter and the CCD sensor have to be well separated from each other to provide a baseline for accurate triangulation. This causes occlusion and anisotropic variation in brightness near sharp depth discontinuities, producing systematic errors.

To explore these errors experimentally we built a precision-milled calibration object. The object consists of bulls-eye shaped grooves of different depths, as seen in Figure 3 left. The large bulls-eyes are 37mm in diameter, and they range linearly in depth from 4.6mm to 1.0mm, providing sharp edges at all orientations with different magnitudes.

De-noised depth maps acquired from the calibration object with our two scanners clearly exhibit similar systematic artifacts as shown in Figure 3. On rising edges (depth decreases sharply along the baseline in the sensor-to-emitter




Figure 3: On the left is a photograph of our bulls-eye calibration object. On the right, we show cross-sections, taken along the sensor-emitter axis, for three bulls-eyes of different edge depths (4.7mm, 3.5mm and 2.3mm). A greater depth discontinuity causes a sharper overshoot, makes the stair-step artifact more prominent, and spreads both of them out over a wider area.

direction), we see a characteristic *stair-step* artifact, with an extraneous bump on the edge. On falling edges, we see an *overshoot*, a small bump occurring just before the edge itself.

A gradient image of the bulls-eye depth map captured from the calibration object is seen in Figure 6b. The errors are most prominent at edges perpendicular to the baseline, decreasing roughly linearly with the angle over a range of about 45° . The greater the depth disparity at the edge, the more significant the systematic errors and the wider the region which they affect (Figure 3 right).

The pattern of systematic error observed with a Cyberware scanner by Curless and Levoy [CL95] had overshoot artifacts on *rising* edges (see their Figure 2c), which surprisingly is inconsistent with our experiments. This may be related to the fact that the Cyberware scanner moved both the laser and camera together, rather than rotating the laser as the Minolta 910 and the NextEngine do.

To explain the errors, we assume that the laser light beam has a Gaussian intensity profile, and that the surface is perfectly Lambertian. At a rising edge (Figure 4, left) the image of the laser stripe is split into two parts. If the peak is detected somewhere between the two parts, triangulation with the erroneous peak places an extra stair-step on the rising edge. On the right in Figure 4, the surface at the falling edge faces the emitter, and is brightly illuminated. We believe that this causes the peak to shift towards the emitter, producing the overshoot error.

Developing a quantitative, computational model of this error process from first principles seems to be far from straightforward. It should be possible to build a data-driven model of the error generation process for a specific scan-




Figure 4: This 1D illustration of systematic error shows the two scenarios in which the error occurs. In both figures we see the laser emitter and the sensor that captures the reflected intensity from an illuminated surface. The figure on the left shows a rising edge. The light emitted from the laser appears to be split, when viewed from the sensor, producing a bi-modal intensity profile in the sensor image, which leads to the stair-step error. The right image illustrates a falling edge. Because of the illumination discontinuity, the reflected intensity profile is skewed, resulting in an overshoot error.

ner, given enough ground-truth data, Unfortunately, our data is “ground truth” in only a limited sense, because although we have a precise model of the shape of the calibration object, we do not know the registration of that shape with the collected data. Every registration we computed of the precise synthetic model with the data was clearly offset slightly along the dimension which suffered from the systematic errors.

4. Removing Systematic Errors

We remove the systematic error by combining several depth maps, captured by scanning the object in different orientations. Our approach (see Figure 5) requires as input a set of acquired depth maps at several orientations $\{h_{\theta_1}, \dots, h_{\theta_n}\}$, differing by a rotation around the z -axis of the scanner (the depth direction in each depth map). We use the term “depth map” rather than “scan” since each input h_{θ_i} is itself expected to be the result of combining several raw scans by averaging or super-resolution. Each h_{θ_i} is used to compute weights w_{θ_i} at each pixel. The weighted input depth maps $\{h_{\theta_1}, \dots, h_{\theta_n}\}$ are then registered and merged to form a single output depth map h_f .

4.1. Weight Computation

Based on the observations in Section 4, we devised a weighting scheme that down-weights the regions of each depth map




Figure 5: Input scans per orientation are combined to produce $\{h_{\theta_1}, \dots, h_{\theta_n}\}$. Weight w_i is computed for each h_{θ_i} . The combined depth maps and weights are used to compute the final depth map h_f free of systematic error.

that we expect to contain errors, that is, the points on edges perpendicular to the baseline. The weights are based on the gradient of the depth function. Since the gradient is sensitive to noise, we compute the weights based on the processed depth maps from which random noise was removed, rather than on the noisy raw input (this is also of course faster).

Edge detection: We use the magnitude of the gradient, $|\nabla f(p)|$, as our edge filter kernel where f is one of the h_{θ_i} . Second-order kernels such as the Laplacian are also often used for edge detection, but we prefer the first-order kernel here because it provides the edge orientation and because second-order kernels tend to amplify noise [dMCOT89].

As is standard in feature detection, we smooth $f(p)$ with a Gaussian $G_\sigma(\cdot)$ before computing the gradient, so that $g(p) = |\nabla(G_\sigma * f)(p)|$. To smooth the stair-step artifacts, which increase with the depth discontinuity, we choose the width σ of the Gaussian to be larger on objects with large discontinuities (see Figure 8). In our experiments, a single choice of σ for each entire object was sufficient, in the range 0.5-3.0 depending on the object.

We can compute the gradient magnitude directly on our input depth map, for reasonably flat objects; but for general




Figure 6: Visualizations of the systematic errors of the 4.6mm depth bulls-eye(a). The gradient magnitude (b) is visualized with blue-to-red representing low to high magnitude, and the gradient orientation (c) with blue-to-red representing the angle from sensor-emitter baseline. Notice in image (b) that the erroneous low gradients due to the stair-step errors on the rising edges appear at edges perpendicular to the baseline, as well as high gradients on falling edges due to overshoot.




Figure 7: Left, before correction for the rounded shape of the parrot, the calculated weight is low (blue) everywhere on the sides the body. Middle, subtracting a smoothed version of the depth map from the original, correctly identifies the sharp horizontal edges. Right, a photograph of the model is shown to convey the roundness of the model.

rounded objects, the gradient magnitude is both a function of the high-frequency detail, with which we are concerned, and the orientation of the surface to the scanner. We take the conventional approach of removing the low-frequency component of the input depth map. We smooth the depth map, and then subtract the smoothed depth map from the original. This 2.5D approach correctly preserves the sharp depth discontinuities in the depth map, rather than the sharp edges of the 3D model. Figure 7 shows an example of the results on a rounded object.

Weight function: The weight function depends on both $g(p)$ and the angle $\theta(p)$ between the direction of $\nabla f(p)$ and the baseline. To generate a smooth result along sharp features, for example on a circular edge where the orientation changes, we need to blend the contributions of the inputs taken from the four different orientations. We smoothly blend the different depth maps by applying a standard sigmoid function $s(\cdot)$ to both $\theta(p)$ and $g(p)$ in the weight function. Our weight function, combining the thresholded values, is then:

$$w(p) = 1 - s(g(p)) \cdot s(\theta(p))$$

4.2. Registration and Merging

We do an initial registration of the weighted input depth maps h_{θ_i} with the recently released 4PCS code [AMCO08]. This algorithm uses a random-sampling method to find sets of four points which, when aligned, register the input surfaces together so as to maximize the closely overlapping surface area. This is followed by an application of the Iterated Closest Point (ICP) algorithm [BM92, RL01] to refine the registration. High-quality registration is important for preserving sharp edges.

After registration, we use the weights to combine the input depth maps h_{θ_i} to form a single output depth map. Output depth values are computed simply as a weighted sum of depth values of neighboring points. The weight assigned to an input point q , at an output point p , is determined by the weight $w(q)$ determined in Section 4.1, as well as a Gaussian function of the distance between them, $G_{\sigma'}(\|q - p\|)$. The standard deviation σ' here is the width of one cell in the grid.

$$z(p) = \frac{\sum_{q \in N} z(q) G_{\sigma'}(\|q - p\|) w(q)}{\sum G_{\sigma'}(\|q - p\|) w(q)}$$

5. Data Acquisition and Preprocessing

Our observations of the calibration object in Section 3 demonstrated that the systematic errors are most prominent at edges perpendicular to the baseline, but they are also evident on edges that differ by as much as 45° (see Figure 6b). Therefore we collected data at four orientations: zero, thirty, sixty and ninety degrees.

With both scanners, the Minolta Vivid 910 and the Next-Engine, we took multiple scans at each orientation, which we used to improve the depth maps' resolution to the point at which the systematic error became visible. In both cases, when taking each scan, we introduced an arbitrary displacement by moving the x and y panning knobs of the laser scanner in very small amounts, manually trying to replace the scanner in the same position every time. This creates random small offsets between the grids of the individual scans, which helps remove quantization error.

With the high-end Minolta Vivid 910, both the random noise level and the resolution at which the systematic error appears is very low. To emphasize the scanner error, we took roughly one hundred raw scans of each of our examples at each orientation and we used super-resolution [KMA06] to remove noise and increase the resolution of the depth




Figure 8: Top, the smoothing factor σ is chosen too small, and the horizontal edges display two regions of low weight (blue). Bottom, a larger σ causes the entire region to have low weight. On the right, we show 1D profiles of the smoothed depth maps.

| | Persepolis | | Parrot | | Mayan | | A Bullsseye | | A Bullsseye | | Gear | |
|--------------------------------------|------------|-------|--------|-------|-------|-------|-------------|-------|-------------|-------|-------|-------|
| Minolta or NextEngine | M | M | M | M | N | N | N | N | N | N | N | N |
| Scans per orientations | 100 | 100 | 100 | 102 | 24 | 24 | 24 | 24 | 24 | 24 | 24 | 10 |
| Scanning per orientation (min) | 20-30 | 20-30 | 20-30 | 20-30 | 15-20 | 15-20 | 15-20 | 15-20 | 15-20 | 15-20 | 15-20 | 10-15 |
| Raw scan resolution (mm) | 0.2 | 0.3 | 0.7 | 0.5 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.12 |
| Points per raw scan | 132K | 59K | 9K | 1.5K | 2K | 2K | 2K | 2K | 2K | 2K | 2K | 143K |
| Preprocessing for h_{θ} (min) | 29 | 15 | 28 | 20 | 11 | 11 | 11 | 11 | 11 | 11 | 11 | 9 |
| Resolution of h_f (mm) | 0.07 | 0.1 | 0.2 | 0.17 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.1 | 0.12 |
| Points in h_f | 1257K | 550K | 1844K | 19K | 63K | 63K | 63K | 63K | 63K | 63K | 63K | 184K |
| Computing h_f (min) | 2 | 3-4 | 3-4 | <1 | <1 | <1 | <1 | <1 | <1 | <1 | <1 | <1 |

Table 1: Our experimental results. The time required to combine the depth maps from the different orientations is small in comparison to data collection and pre-processing.

maps at each orientation by a factor of three. The super-resolution algorithm uses an iterative registration and surface reconstruction pipeline, similar to the one recently adopted by [HAW07] for surface reconstruction.

With the inexpensive NextEngine scanner, the laser beam is wider, and the systematic error is clearly evident at the resolution of the scans, once random noise is removed. We used roughly ten to twenty scans per orientation, and averaged them using the iterative registration and surface reconstruction pipeline of the super-resolution algorithm, but without actually attempting to increase the resolution. This removed both random noise and striping artifacts, but greatly emphasized the systematic error, as seen in Figure 9 and Figure 11.

6. Results and Discussion

We begin with the results for the NextEngine scanner, where combining multiple scans to improve the quality of the output proved to be a successful strategy. On the bulls-eye calibration object we can remove the systematic error almost perfectly, as seen in Figure 9. We also scanned a gear from inside of a toy gumball machine, shown in Figure 11. The gear is about 5cm in diameter. As can be seen, the systematic errors in the scans of the gear are very pronounced. Nonetheless they could be removed using the same technique, producing a much improved output depth map.

Super-resolved depth maps taken with the Minolta Vivid 910 are of very high quality except for the systematic error. In the bulls-eye calibration object, Figure 2 shows one bulls-eye with grooves of 4.6 mm deep, represented on a super-resolution grid at .17 mm resolution. Our processing successfully removes the systematic errors. In Figure 12 we show our output depth map from scans of cast of a Mayan hieroglyphic (25cm wide). Here again, the technique was very successful.

Our technique is quite simple, and one might wonder if something even simpler might work as well. In Figure 12c, we show the result of combining the four input orientations without using the weighting of Section 4.1. The system-

atic error remains noticeable, which demonstrates the importance of our weighting scheme. In another experiment, we removed any points at which the surface normal deviated by more than 30 degrees from the scanner’s z -axis during the acquisition phase. Again we combined the trimmed scans without using the weighting scheme. Not only does this remove data that does not have the systematic error, introducing holes, but it also sometimes fails to remove the horizontal part of the stair-step error, for instance near the spear in Figure 10.

Table 1 shows the resolution and timings for our examples. Although the resolutions of the models scanned with the NextEngine are nominally higher than those scanned with the Minolta, not only is the signal-to-noise ratio higher, but high-frequency features seem to be smoothed away.

The technique worked well on surfaces which have sharp edges, deep grooves, or other large features, and also on surfaces with anisotropic detail. Figure 13 shows a close-up of the back of the parrot statue which appears in Figure 7. While the systematic errors incorrectly emphasize the fine vertical details in the feathers in one of the input orientations (13b), the output is dominated by the scans from the other orientations, in which the edges are correct (13a).




Figure 9: Top, a single scan from the NextEngine Scanner. The systematic error is masked by noise. Center, averaging 24 such scans removes most of the noise and reveals the typical pattern of systematic error. At the bottom, merging four scans using our method produces a final depth map that is dramatically better than the input scans.

7. Limitations and future work

We found that the NextEngine had to be positioned very carefully to fully capture the details of the bulls-eye and the gear surface, since the range within which we could capture high frequencies well was rather small. Also, because NextEngine’s scanning speed is slow, capturing multiple nearly

identical scans in order to improve its quality was time consuming. This remains an obstacle to our idea of using large amounts of data in order to improve the depth maps captured with the NextEngine. On the other hand, it is clear that combining multiple NextEngine scans does improve quality, eliminating random noise, striping artifacts, and systematic scanner error.

There are a number of lessons learned from this project that should be applicable to future work. Areas in which our pipeline could be improved include using the weight functions in the registration process, and adaptively choosing the smoothing factor σ during edge detection, perhaps with a framework such as SIFT [Low99].

Similar techniques could be applied in future work to the problem of removing systematic errors at reflectance discontinuities. Finally, while we feel that we have made some contribution towards understanding systematic error in triangulation laser range scanners, we have fallen short of the goal of developing a computational model for its simulation or removal.




Figure 10: A naïve approach to systematic error removal: average all the input scans by first removing points at which the surface normal deviates by more than 30 degrees from the scanner z-axis. Although large portions of the scans were removed, systematic errors remain such as the stair-step artifact on the right side of the spear. A good amount of data was lost, such as the holes at the left side of the spear. Similarly, the nose of the soldier should have a smooth boundary instead of a sharp one. See Figure 1 for comparison.

8. Acknowledgments

This work was supported by NSF grants CCF-0635250 and CCF-0401601.

References

[AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust pairwise surface registration. *ACM Transactions on Graphics (TOG) (Proceedings of SIGGRAPH)* 27, 3 (2008), 1–10.

[BLS92] BUZINSKI M., LEVINE A., STEVENSON W. H.: Performance characteristics of range sensors utilizing optical triangulation. In *NAECON '92: Proceedings of the IEEE 1992 National Aerospace and Electronics Conference* (May 1992), pp. 1230–1236 vol.3.

[BM92] BESL P., MCKAY N.: A method for registration of 3D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 2 (Feb 1992), 239–256.

[BR86] BLAIS F., RIOUX M.: Real-time numerical peak detector. *Signal Processing* 11, 2 (1986), 145–155.

[CL95] CURLESS B., LEVOY M.: Better optical triangulation through spacetime analysis. In *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision* (1995), pp. 987–994.

[dMCOT89] DE MICHELI E., CAPRILE B., OTTONELLO P., TORRE V.: Localization and noise in edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 10 (1989), 1106–1117.

[DTB06] DIEBEL J. R., THRUN S., BRÜNIG M.: A bayesian method for probable surface reconstruction and decimation. *ACM Transactions on Graphics (TOG)* 25, 1 (2006), 39–59.

[FN96] FISHER R. B., NAIDU D. K.: A comparison of algorithms for subpixel peak detection. In *Proc. 1991 British Machine Vision Association Conf* (1996), pp. 217–225.

[FSCP04] FOREST J., SALVI J., CABRUJA E., POUS C.: Laser stripe peak detector for 3D scanners. a FIR filter approach. In *ICPR '04: Proceedings of the 17th International Conference on Pattern Recognition* (2004), vol. 3, pp. 646–649.

[HAW07] HUANG Q.-X., ADAMS B., WAND M.: Bayesian surface reconstruction via iterative scan alignment to an optimized prototype. In *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing* (2007), pp. 213–223.

[KGC91] KANADE T., GRUSS A., CARLEY L.: A very fast vlsi rangefinder. In *ICRA '91: Proceedings of the 1991 IEEE International Conference on Robotics and Automation* (April 1991), vol. 2, pp. 1322–1329.

[KMA06] KIL Y. J., MEDEROS B., AMENTA N.: Laser scanner super-resolution. In *Eurographics Symposium on Point-Based Graphics* (July 2006), pp. 9–16.

[Low99] LOWE D. G.: Object recognition from local scale-invariant features. In *ICCV '99: Proceedings of the Seventh IEEE International Conference on Computer Vision* (1999), p. 1150.

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHY R.: Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics (TOG) (Proceedings of SIGGRAPH)* 24, 3 (2005), 536–543.

[RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. In *3DIM '01: Proceedings of the Third International Conference on 3D Digital Imaging and Modeling* (2001), pp. 145–152.

[RTF*04] RASKAR R., TAN K.-H., FERIS R., YU J., TURK M.: Non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. *ACM Transactions on Graphics (TOG) (Proceedings of SIGGRAPH)* 23, 3 (2004), 679–688.

[SLPA90] SOUCY M., LAURENDEAU D., POUSSART D., AUCLAIR F.: Behaviour of the center of gravity of a reflected gaussian laser spot near a surface reflectance discontinuity. *Industrial Metrology* 1, 3 (1990), 261–274.




Figure 11: An input scan of a gear, taken with the NextEngine scanner (left). Averaging ten scans (mid-left) removes general noise, but very large systematic errors become evident. Another de-noised depth map (mid-right), scanned in a different orientation shows equally large systematic errors, but in different regions. Finally, combining the depth maps captured at the four different orientations (right) produces a high-quality output depth map, dramatically better than any of the input scans.



Figure 12: Systematic error removal from scans of a cast of a Mayan hieroglyphic. The output depth map is in the center. On either side, close-ups of selected regions, with cross sections highlighted and shown beneath. (a) shows the super-resolved depth map from scans taken at zero degrees, with systematic errors near the deep horizontal grooves. (b) is the super-resolved depth map taken at ninety degrees, with errors near vertical edges. The depth map created by equally averaging all four input orientations with equal weight shows in (c). Finally (d) shows our result. While the averaging method in (c) does reduce the systematic errors, it clearly fails to eliminate them, indicating that our weighting scheme is necessary.




Figure 13: Fine feather features of the parrot model from Figure 7. The super-resolution scan (a) taken at zero degrees. There are systematic errors in the vertical direction but the mostly vertical fine details in the feathers are captured correctly. In the scan taken at ninety degrees (b), the systematic errors, in particular the overshoot, serves to emphasize the texture. The combined final output (c) is similar to a very high-resolution scan (d), taken by focusing the laser scanner at the small region shown.