

Laser Scanner Super-resolution

Yong Joo Kil¹ and Boris Mederos² and Nina Amenta¹

¹Department of Computer Science, University of California at Davis

²Instituto Nacional de Matematica Pura e Aplicada - IMPA

Abstract

We give a method for improving the resolution of surfaces captured with a laser range scanner by combining many very similar scans. This idea is an application of the 2D image processing technique known as super-resolution. The input lower-resolution scans are each randomly shifted, so that each one contributes slightly different information to the final model. Noise is reduced by averaging the input scans.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Surface Acquisition

Laser range scanners are used to capture the geometry of three-dimensional objects. They are used for reverse engineering manufactured objects, and for digitizing objects of scientific, artistic or historical importance for archiving and analysis. Like all physical measurement devices, laser range scanners have limitations, including noise and limits on resolution. We describe a method for acquiring better models using a commercial laser range scanner by taking many almost, but not quite, identical scans and combining them in software. For instance, in Figure 2, we combined 100 nearly identical scans like the one on the left to create the cleaner and more detailed surface in the center.

Our technique is a variant of a two-dimensional image processing technique called *super-resolution*, which takes many nearly identical low-resolution input images and combines them to produce a higher-resolution output image. Classical super-resolution is based on the assumption that each pixel in the low-resolution input images is produced from an original continuous image by deforming (possibly just translating), blurring and then sampling. If the blur function is a perfect low-pass filter, then all high-resolution information is irrevocably lost. But if not, then the resulting aliasing in the low-resolution images contains information, and a higher-resolution output can be recovered from enough slightly displaced input images.

In its simplest form, this idea is obvious. Consider sampling a signal at some rate less than its Nyquist frequency (the rate at which all features can be captured without aliasing). One set of samples is not sufficient to reconstruct the signal. But if we are given many samples, randomly offset

from each other, and we can register them correctly, then we can produce an accurate reconstruction. See Figure 1. We apply this simple idea to laser range scanning very directly.




Figure 1: Left: If a scan samples the surface too sparsely, detail is lost. The Nyquist theorem implies that we can only expect to capture details of size at least x if we scan with a sample spacing of $x/2$. Right: When we combine several scans (different shades), we can recover much more detail.

Combining many scans is also useful for removing noise. Noise in the depth images from laser range scans comes from several sources, including quantization and noise in the video imaging system, laser speckle (caused by random reinforcement of the coherent light of the laser reflected from a rough surface), systematic error in peak detection (caused by surface curvature and color), and the instability of the computation of point locations by triangulation. Random noise can be essentially eliminated, but super-resolution does nothing for systematic scanner errors.

Obviously, it is much more time consuming to take a hundred scans than it is to take one. In high-value situations,




Figure 2: On the left, one scan of the the parrot statue, with a sample spacing of about 1mm. Center, we combine 100 nearly identical such scans to produce the surface in the center, produced on a grid with sample spacing of about 0.3mm. Notice the noise reduction and the improvement in the detail, for instance in the face, neck and wing feathers. On the right, a photograph of the parrot statue.

however, for instance when an expert has flown thousands of miles to scan a fossil or a coin in a museum, it would be worth it to spend an hour, or even a few hours, scanning in order to get a better model.

Our technique also gives an alternative approach for capturing high-resolution models of large objects. Instead of capturing many high-resolution scans, each covering a small area, and then registering them, we can capture many (slightly shifted) low-resolution scans covering a larger area, and then merge them to produce a single large high-resolution surface patch. This technique has two advantages. First, it avoids having to merge many small scans, which is difficult to do accurately. Second, it means that the scanner can be kept farther away from the object, and moved less, which might be necessary or desirable.

Algorithm Overview: We scan a model many times, from similar but randomly perturbed viewpoints. We get an initial registration of the scans to each other, and we use the registered scans to reconstruct the super-resolved depth-map. We get the depth value at every point of a higher-resolution grid by simply averaging the z-values of nearby points from the input scans. We then re-register each scan to the super-resolved depth-map, and iterate the reconstruction and registration steps several times. Finally we output the super-resolved depth-map.

1. Related Work

Laser range scanners: Producing better models from laser range data is a topic of on-going interest in computer graphics and computer vision. Curless and Levoy [CL95] give an excellent description of the scanning process. With a scanner such as our Minolta Vivid 910, a vertical stripe of laser light is moved across the object surface, and captured by a video camera. Along each horizontal scan line of the video frame, the brightest spot is taken to be the point at which the laser stripe "hits" the surface. This brightness peak is detected at sub-pixel resolution. The relative positions of the laser and the video camera are used to find the three-dimensional coordinates of the brightest spot by triangulation. So, the x-coordinate of each point in the output depth image is determined by the position of the laser stripe for a particular video frame, the y-coordinate corresponds to a raster line in the video frame, and the depth value is computed from the brightness peak detected along the raster line in the video frame.

This imaging process is more complex than that of a camera, and it introduces systematic artifacts. Curless and Levoy removed many of these, and reduced noise, by analyzing multiple frames of the video stream when detecting each peak point. Their technique is intended to be part of the processing within the scanner. We concentrate on improving noise and resolution at the user level, taking the output data from the device and trying to improve it by post-processing.

Recently Nehab et al. [NRDR05] combined surface data from a temporal stereo triangulation scanner with normals captured by photometric stereo, to improve the resolution and reduce the noise in captured surface models. Similarly, Diebel and Thrun [DT05] improve models captured with a time-of-flight scanner using digital photographs. Rather than combining different data types, we explore combining lots of scans.

Our reconstruction method is related to Curless and Levoy's VRIP algorithm [CL96], with one crucial difference. They treat each input scan as a triangle mesh, with the intention of getting as much coverage as possible. Interpolating the scans by triangles is a form of low-pass filtering, and high-resolution detail is lost, or at least greatly obscured. Instead, we treat each scan as a set of points, making it much easier to reconstruct a higher-resolution surface.

Super-resolution: Our method is derived from two-dimensional super-resolution algorithms. See [PPK03] for a recent survey. Most algorithms use two basic steps, sub-pixel registration and reconstruction. Neither step translates completely straightforwardly into our context. Reconstruction is usually the more complex step, and is a large focus of our research. It is based on the assumption that, given the correct registration, each low-resolution input image is formed from some "true" high-resolution image by a known process: first the "true" image is translated slightly (or displaced in some other way), then it is smoothed with a blurring kernel (e.g. a box filter), then it is sampled at the lower resolution, and finally random noise is introduced. Elad and Feuer [EF97] argue that most methods can be represented by the following matrix equation

$$Y = HX + E$$

where Y is a vector containing all the low-resolution images, H is a known matrix that contains the displacement, blurring and sampling steps of the image formation process described above, X is the desired high-resolution image, and E is an unknown vector of noise terms. This basic model can be enhanced with priors on the error terms and/or regularization terms on the high-resolution pixels to encourage smoothness. A least-squares solution, minimizing $E^t E$ is formulated in the standard way, and the problem comes down to solving a huge sparse linear system to recover X , which in practice are solved using an iterative methods.

An ideal super-resolution algorithm for laser range data would be based on a matrix H which correctly models the triangulation laser range scanner. This seems difficult since the peak-detection process is not a linear operation. Instead, we assume an admittedly incorrect but linear process: translation and down-sampling, with no blur operation. The main advantage of this assumption is that it leads to a very simple and efficient reconstruction algorithm. The alternative assumption, that nearby depth values are averaged by the peak-detection process, is also clearly incorrect (and so far in our research, has not produced better results).

Our simple reconstruction algorithm is based on an argument by Elad and Hel-Or [EH001], in which they point out that in very easy cases, the correct solution for X can be found by interpolating the values in Y onto the grid of values for X . We discuss their argument as it applies to our case in Section 5. It is very similar to the splatting-interpolation process for surface reconstruction from laser range scans used in VRIP [CL96], which Curless similarly showed it minimized the least-squares error in his thesis [Cur97].

The large-scale structure of our algorithm is inspired by the super-resolution algorithm of Cheeseman et al. [CKK*96], in which after producing a super-resolution image, each input image is then re-registered to the super-resolution output, and the process is iterated. This gives us the opportunity to improve the registration using high-resolution features which are not detectable in any of the low-resolution inputs.

2. Data Acquisition

We take around 100 scans of each view of the object. To create random displacements between scans, we move the entire scanner on its tripod before each scan, by nudging its x and y panning knobs ($\pm 5^\circ$). This is a laborious process, requiring 30-60 minutes of scanning for each view. It would also be possible to move the object slightly before each scan, instead of moving the scanner, possibly by using an automatic turntable.

For the Mayan hieroglyphic model in Figure 10, we also rotated the tablet after every ten scans. This reduced "tiling" artifacts caused by interaction between the low-resolution and high-resolution grids, as discussed in Section 6.

We used the Geomagic Studio [Geo03] software to clean up the scans, removing pieces of surrounding objects and supports. This is a normal part of the 3D scanning pipeline. The scans could be manipulated altogether as a group, so this stage is not more time consuming than usual. We also used Geomagic's registration tool to produce an initial registration of the scans, choosing an arbitrary scan and registering all the others to it. This registration is improved by the later super-resolution processing.

Our goal is to capture enough three-dimensional points so that at least a few scan points contribute to each high-resolution output point. To get a ball-park figure, assume that we want to improve the resolution by an integer factor m (such as four), and assume that the displacements are simply translations of the low-resolution grid with respect to the high-resolution grid. Then if we consider an arbitrary $m \times m$ square of high-resolution cells, each low-resolution scan contributes one point to the square, which lands randomly in one of the cells. A scan misses each specific cell with probability $(1 - 1/m^2)$, so the probability that any of the m^2 cells end up with fewer than, say, five points after we take k scans is $\leq m^2(1 - 1/m^2)^{(k-5)}$. Then k , the number

of scans needed, grows as $O(m^2 \lg m)$, and, to give a specific example, if $m = 4$ and we want the failure probability to be < 0.1 , we need about 150 scans. In fact we usually use fewer. While increasing the number of scans improves the quality of the surface - see Figure 3 - it seems indeed to do so only slowly.




Figure 3: Super-resolution reconstruction using only 30 input scans at the left and increasing to 140 at the right. Noise is reduced dramatically at the beginning but more slowly at the end. Surfaces were reconstructed from subsets which were pre-registered using all 140 scans.

3. Reconstruction

Our reconstruction algorithm is very simple. We produce an output depth-map as a grid of function values, each computed as a weighted average of the contributions of nearby points from the input scans.

We first establish a coordinate system in which the z -direction (depth) is the direction towards the scanner, by adopting the coordinate system of an arbitrarily chosen input scan. We create a grid of the desired resolution in the $x - y$ plane of that coordinate system, and locate each of the scanned input points in the grid, based on the current registration. Then for each high-resolution cell with center q we create an output depth value. We use a two-dimensional Gaussian kernel centered on q to assign weights $w(q, r_i)$ to the input points r_i in cells within the surrounding 5×5 neighborhood N , based on their $x - y$ positions. Then we take the weighted average of their depth values.

$$z(q) = \frac{\sum_{r_i \in N} z(r_i) \cdot w(r_i, q)}{\sum_{r_i \in N} w(r_i, q)}$$

The width h of the Gaussian kernel $e^{-d(r_i - q)^2 / h^2}$ is set to the sample-spacing of the grid, so that points outside of the 5×5 neighborhood have negligible weight. The sum is dominated by the points within the 3×3 neighborhood, and when $m > 3$ (as in all our experiments) this neighborhood can contain at most one point from each scan.

We experimented with using the median (which is more robust to outliers) rather than the weighted mean, and we




Figure 4: A thin strip of the super-resolved surface, and the nearby sample points from the input scans. The input is very noisy, but the points are densely and randomly distributed near the surface with few outliers, so the average gives an accurate representation of the surface.

did not see any noticeable difference in the output. As seen in Figure 4, the noise seems fairly evenly distributed near the surface, and outliers are not much of a problem, so a least-squares model is probably best.

The Gaussian weights provide some smoothing as part of the interpolation, but the resulting surface still shows some noise, as in Figure 5. We follow the reconstruction by applying a bilateral filter [TM98], removing some noise while retaining sharp features.




Figure 5: A close-up of the reconstruction of the parrot-head model before bilateral filtering (left) and after (right). We re-scanned the head of the parrot, with a sample spacing of about .4mm in the input scans, and reconstructed the super-resolution surface from 146 scans. The entire surface appears in Figure 11.

4. Registration

We iterate reconstruction of the surface with re-registration of each scan to the high-resolution surface. We use the following variant of the Iterated Closest Point (ICP) algorithm [BM92, RL01] to register a scan to the surface. First we subsample the triangles of the high-resolution surface, distributing samples in the interior of each triangle. Then for each point in the scan, we find its closest point in the dense set of samples on the surface. Points near boundaries, either on the scan or on the surface, are ignored. We use Horn's algorithm [Hor87] to minimize the mean-squared distance between the point pairs, and iterate until convergence. We

see an example of the improvement produced by registering each scan to the high-resolution surface in Figure 6.




Figure 6: Left, a close up of the initial super-resolution mesh, for which the scans are only registered to each other. Right, after one iteration in which the scans are registered to the high-resolution surface. Notice that the "tiling" artifact typical of mis-registration is reduced. We see the most dramatic improvement in this first step.

5. Least-Squares Approximation

As noted in Section 1, averaging minimizes the error in a least-squared sense, which corresponds to a maximum-likelihood estimate of the surface, assuming the errors in the z -direction are Gaussian noise. Here we review the argument of Elad and Hel-Or [EHO01] which makes this assertion precise under some assumptions that are not strictly met in practice, but which are a good rough approximation. Our set-up is somewhat different from theirs in that we assume there is no blur operation. This allows us to drop one of their assumptions that the deformations used to form their input images are all translations in the $x - y$ plane, but we can allow rotations as well.

The raw input scans are represented as column vectors $\{Y_1, \dots, Y_L\}$ with dimensions $n \times 1$, where $n = \text{width} \times \text{height}$. The high-resolution depth-map is represented as a column vector X of dimension $N \times 1$, where $N = \lceil \text{width} \cdot m \rceil \times \lceil \text{height} \cdot m \rceil$.

Based on our assumption that there is no blurring and assuming that the range scans are obtained by random rigid motions, the creation of the low resolution scans Y_k , $k = 1, \dots, L$ from an unknown high-resolution depth-map X could be explained by the following linear model:

$$\{Y_k = D_k F_k X + E_k\}_{k=1, \dots, L} \quad (1)$$

The matrices D_k , F_k , E_k represent a known decimation, a known displacement, and an unknown Gaussian error vector, respectively. What we want to find is a high resolution X such that $E = \sum_{k=1}^L E_k^t E_k$ is minimized. Expressing E as a function of X and setting the derivative of $d(E)/dX = 0$, we get a linear system

$$RX = P \quad (2)$$

where $R = \sum_{k=1}^L F_k^t D_k^t D_k F_k$ and $P = \sum_{k=1}^L F_k^t D_k^t Y_k$. Generally, in super-resolution algorithms, the linear system is solved via an iterative method such as steepest descent.

Following Elad and Hel-Or, we make the following additional simplifying assumptions. Assume that both D_k and F_k

Model	Dimension (meters) (width × height)		Resolution (mm)	# points (per scan)	# of Scans
N.Y. Subway Token	0.02	0.02	0.15	24k	117
Parrot Head	0.026	0.023	0.4	3k	146
Parrot Whole (single view)	0.06	0.2	1.5	10k	100
Mayan Hieroglyphic	0.24	0.18	1.0	47k	90

Table 1: Statistics on all the scanned input data sets.

are $0 - 1$ matrices, where D_k represents decimation exactly by an integer factor m , and each F_k is a permutation approximating the displacement which it represents. Then R turns out to have a simple form: it is an integer diagonal matrix, where the (i, i) th entry is the number of samples placed into the i th cell of the high-resolution grid. The column vector P is simple as well: the i th entry contains the sum of all the z values of samples that fall into the i th high-resolution cell. So under these assumptions, averaging the z values of the points that fall into the i th cell is the correct solution of the linear system, and gives the least-squares estimate of the z value of the surface at that point.

6. Experiments

We have tried this method on models of various sizes, with the statistics summarized in Table 1. For each, the resolution of the output mesh was either 3.7 or 4 times that of the input scans. We do five iterations of reconstruction and registration for each super-resolved scan (three would probably have been sufficient), and each iteration takes between one and three minutes, depending of the size of the input model and the desired output resolution. Actual scanning time takes anywhere between 30-60 minutes for 90-150 scans, and cleanup and initial registration using Geomagic takes another 30 minutes or so.

We took six views of the parrot model, producing six super-resolved scans. As we can see in Figure 2, our super-resolved scans have artifacts that are derived from systematic scanner errors which occur in all input scans. In particular, because of "edge curl" (for instance at the wing tips) the boundaries of each super-resolved scan had to be trimmed. We used the standard clean-up and merging tools in Geomagic to merge the six super-resolved scans into a single model, which can be seen in Figure 9. In this example it is clear that both noise removal and true super-resolution - the detection of detail invisible in each individual scan - is achieved.

Each super-resolved scan in the parrot model was composed of 100 input scans, each taken with a random $x - y$ pan. Registering and merging the low-resolution point sets produces the periodic sampling patterns on the left in Figure 7; the period matches that of the low-resolution input grid, and is larger than the sample spacing in the output high-resolution grid. This shows up in the final surface as noticeable aliasing on the output surface, which can be seen in the close-up in Figure 6. As registration improves, this aliasing




Figure 7: *Left: the sampling pattern when we only do random x-y shifts. Right: the sampling pattern when we also rotate the model (increments of 10° degrees for every 10 scans, 90 scans total). Notice that this reduces the tiling artifact.*

diminishes but it does not disappear. It is possible that the scans form clumps, and that registration is excellent within a clump but loser between clumps.

To avoid this undesirable artifact, we can rotate the model while taking the scans. We did this with the Mayan hieroglyphic shown in Figure 10, producing the much more desirable sampling pattern on the right in Figure 7. This did indeed eliminate the grid aliasing artifact, at the cost of slightly worse registration as measured by RMS error. While the results overall in this experiment were good, another artifact is visible near the deep grooves in the surface: ridges inside the grooves, sometimes (like in the wrinkles near the bird’s smile) obscuring the groove itself. We believe this occurs because of the peak detection errors on the low-resolution CDC image in the scanner. This highlights a problem with using super-resolution alone, without better processing within the scanner, e.g. Curless and Levoy [CL95]; if there are systematic scanning artifacts, super-resolution succeeds only in improving the resolution of the artifacts.

We tested the method on some very small objects, to show that we can improve the effective resolution of the scanner. The smallest object we scanned was a cast of a New York subway token, shown in Figure 8. The subway token was scanned at the highest resolution possible, with a sample spacing of 0.15mm. This is in fact a higher resolution than the specified accuracy of the scanner (± 0.22 mm in x -direction and ± 0.16 mm in y -direction). The super-resolved model clearly achieves very good output quality at this high resolution, effectively capturing the object completely. To compare the contribution of noise reduction and resolution improvement to this result, we tried processing a single scan to reduce noise, by subdividing and then smoothing. This was surprisingly effective (Figure 8, lower left), but clearly not as good as super-resolution. We also tried taking 100 scans without nudging the scanner before each scan, to see the effect of the small random shifts. The results were again quite good (Figure 8, lower right), but artifacts of the low-resolution grid are clearly visible, and the full super-resolution method is again significantly better. For very small objects like coins, it is clear that taking many scans reduce noise greatly, and that small displacements between the scans and the full super-resolution process give the best results.

At a lower input resolution, we get excellent super-resolution results. A high-resolution scan of the parrot’s head, with an input sample spacing of about 0.4mm, can be seen in Figure 11. In this case details completely invisible in the input scans really are revealed in the super-resolved output, especially the differences in texture.

The high quality renderings in Figures 2, 9, 10 were done using Autodesk Maya, while the blue models are screenshots from Geomagic Studio. All models are triangulated.

7. Discussion

We see several avenues for further research based on this method. Our reconstruction algorithm uses the simplest approach to super-resolution. Most image processing algorithms involve more complicated formulations. We tried one such approach, based on the super-resolution algorithm of Irani and Peleg [IP91] on our Parrot Head example, and did not see a noticeable improvement. Possibly other super-resolution algorithms, for instance [FREM03] or [KJH01], could improve the results using better regularization and filtering terms, and might give good results using less input data, as they seem to do in image processing.

Our parrot model is constructed from six super-resolved scans. Instead, it might be possible to use a turntable and use several hundred input scans from different directions to produce a super-resolved cylindrical scan. This would require a true 3D reconstruction method, rather than the simple 2.5D processing scheme we used here. Also, super-resolution requires large quantities of data, and the limiting factor in our prototype is memory. Organizing the computation to use memory efficiently would make it feasible for large scanning projects. A system similar to Curless and Levoy’s VRIP method [CL96] might be appropriate.

Finally, we are eager to try super-resolution with time-of-flight scanners. Time-of-flight scanners suffer from noise in the z -direction, and they fit the model of pure point-sampling very well.

8. Acknowledgments

We gratefully acknowledge the cheerful assistance of U.C. Davis undergraduate Research Assistant Kelcey Chan, who helped us with the data acquisition and model construction. We thank the reviewers for several helpful suggestions.

Work on this project was supported by NSF through grant CCF-0331736 and also through CAREER award CCF-0093378. Dr. Mederos was also supported by the Brazilian National Council of Technological and Scientific Development (CNPq).




Figure 8: (a) One scan. (b) Final super-resolved surface from 100 scans. (c) Photo of the object (a plaster cast of a subway token). The bottom row shows some results of other kinds of processing, to evaluate the importance of the various steps of the algorithm. (d) One scan, bilinearly interpolated onto the finer grid and smoothed. Detail is missing. (e) The entire algorithm except for the final bilateral filtering step. The noise removed by the filtering seems to be residual registration error, which perhaps could be improved. (f) Just averaging 100 scans taken without moving the scanner, using the same Gaussian kernel. Noise is decreased, but there is aliasing from the lower-resolution grid obscuring detail visible in (b).

References

- [BM92] BESL P., MCKAY N.: A method for registration of 3d shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 14, 2 (Feb 1992).
- [CKK*96] CHEESEMAN P., KANEFSKY B., KRAFT R., STUTZ J., HANSON R.: Super-resolved surface reconstruction from multiple images. In *Maximum Entropy and Bayesian Methods*. Kluwer Academic Publishers, 1996, pp. 293–308.
- [CL95] CURLESS B., LEVOY M.: Better optical triangulation through spacetime analysis. In *5th International Conference on Computer Vision* (1995), pp. 987–994.
- [CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. *Computer Graphics* 30, Annual Conference Series (1996), 303–312.
- [Cur97] CURLESS B. L.: *New Methods for Surface Reconstruction from Range Images*. Tech. Rep. CSL-TR-97-733, Stanford Computer Science, 1997.
- [DT05] DIEBEL J., THRUN S.: An application of markov random fields to range sensing. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)* (Cambridge, MA, 2005), MIT Press.
- [EF97] ELAD M., FEUER A.: Restoration of single super-resolution image from several blurred, noisy and down-sampled measured images. *IEEE Trans. on Image Processing* (1997), 1646–1658.
- [EHO01] ELAD M., HEL-OR Y.: A fast super-resolution reconstruction algorithm for pure translational motion and common space invariant blur. *IEEE Transactions on Image Processing* 10, 8 (2001), 1187–1193.
- [FREM03] FARSIU S., ROBINSON D., ELAD M., MILANFAR P.: Robust shift and add approach to super-resolution. In *SPIE Conf. on Applications of Digital Signal and Image Processing* (2003).
- [Geo03] GEOMAGIC I.: Studio 6.0, 2003.
- [Hor87] HORN B.: Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America* 4, 4 (1987), 629–642.
- [HT84] HUANG T., TSAY R.: Multiple frame image restoration and registration. In *Advances in Computer Vision and Image Processing*, Huang T., (Ed.). 1984, pp. 317–339.
- [IP91] IRANI M., PELEG S.: Improving resolution by image registration. *CVGIP: Graph. Models Image Process.* 53, 3 (1991), 231–239.
- [KJH01] KIM H., JANG J.-H., HONG K.-S.: Edge-enhancing super-resolution using anisotropic diffusion. In *IEEE Conference on Image Processing* (2001), vol. 3, pp. 130–133.
- [NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHY R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics (SIGGRAPH 2005)* (2005), 536–543.
- [PPK03] PARK S. C., PARK M. K., KANG M. G.: Super-resolution image reconstruction: a technical overview. *IEEE Signal Processing Magazine* (2003), 21–36.
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the icp algorithm. In *International Conference on 3D Digital Imaging and Modeling (3DIM)* (2001).
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision* (Washington, DC, USA, 1998), IEEE Computer Society, p. 839.



Figure 9: Six super-resolved scans, merged to form a complete model. The original scan data (Figure 2, left) was quite noisy. The super-resolved scans include scanner artifacts, particularly curling at the edges, which had to be trimmed interactively before merging.



Figure 10: On the left, a single scan of a cast of a Mayan hieroglyphic, with a sample spacing of about 1mm. Center, 90 scans are combined to make a super-resolution surface. Notice the improvement on the eyes, the cross-hatching, and the area in front of the bird's face. Right, a photograph of the cast.




Figure 11: On the left, a single close-up scan of the head of the parrot statue, with sample spacing about 0.4mm. Center, 146 scans are combined to make a super-resolution surface. Notice the feather texture on the face, which was invisible in the single scan. Right, a photograph of the statue.