Combining Laser Scans

Yong Joo Kil *
University of California, Davis

Boris Mederos †
IMPA, Rio de Janairo, Brazil

Nina Amenta ‡
University of California, Davis

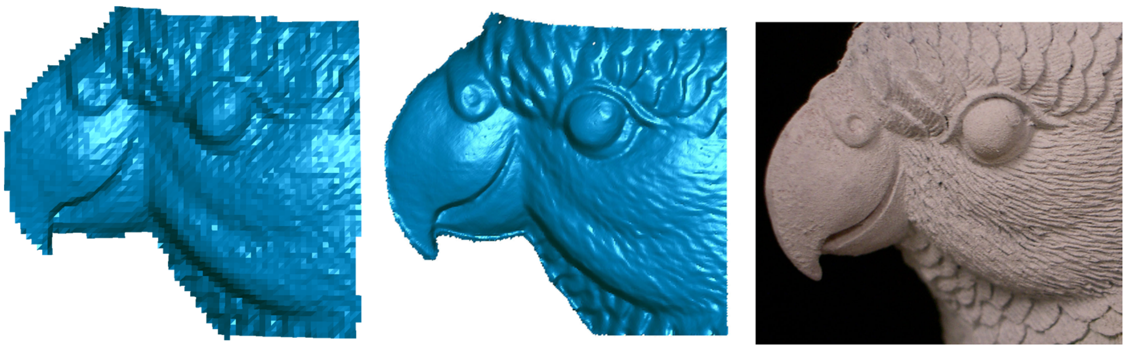

Figure 1: Left, a patch of a single surface captured by a laser range scanner at about .4mm/sample. Center, we integrate many of these surfaces to create a single high-quality surface on a higher-resolution grid. Right, a photograph of the object. Notice the improved detail on the eye and the nostril, the feather texture on the face (invisible in the low-resolution scan), and the overall noise reduction.

We describe method for combining many copies of a surface captured with a laser range scanner to produce a single high-quality surface. The method is inspired by the 2D image processing technique known as *super-resolution* [Park et al. 2003]. The input scans are randomly shifted, so that each scan contains slightly different information, which contributes to the final model. Noise is reduced by averaging.

In image processing, super-resolution is based on modeling the intensity of each low-resolution pixel as a function of the unknown intensities of some set of high-resolution pixels (determined by a registration of the input images), and then solving for the optimal high-resolution image. This reconstruction can be iterated with registration until convergence. For laser range scans, the low-resolution data are more like point samples from a depth map than averages of depth, so we model each scan simply as an array of point samples with noise. We iterate registration of the scans with a reconstruction process which simply averages nearby input points to produce each output depth value.

Implementation

We implemented a 2.5 dimensional version of this idea, which produces a depth image. To create Figure 1, we took 146 input scans, each containing about 4.5K 3D points, using a Minolta Vivid 910 Laser Scanner

We stochastically perturbed the scans by slightly moving the scanner on its tripod in the x and y axis by an arbitrary amount, by hand, before each scan. Capture took less than an hour.

To initialize the processing, we choose one of the of the N raw input scans $\{R_1, \ldots, R_N\}$ arbitrarily to determine a global x, y, z coordinate system, with the z direction pointing towards the scanner. We register the other scans to it in 3D using the Iterative Closest Point (ICP) algorithm.

We then reconstruct the high-resolution depth-map as the z-function, on a grid Q in the (x, y)-plane, with grid spacing h. We set $h = h_r/m$, where h_r is roughly the spacing of the sampling grid of the input scans, and m is the difference in grid spacing, 4 in our example. Assuming the perturbations of the input scans are random, with $N = \Omega(m^2 \lg m)$ scans there is at least one raw sample r_i per grid cell with high probability.

We reconstruct the z value in each cell by averaging the samples within a small neighborhood, using a Gaussian kernel. This reduces noise in the z-direction and compensates for the uneven distribution of samples. For each grid location $q \in Q$, we take a weighted average of nearby raw samples r_i:

$$z(q) = \frac{1}{k} \sum_{r_i \in N(q)} z(r_i) \cdot w(r_i, q)$$

Here $N(q)$ the 25-cell neighborhood around q, $k = \sum_{r_i \in N(q)} w(r_i, q)$, and the weight $w(r_i, q) = e^{-\|r_i - q\|^2/h^2}$, where $\|\cdot\|$ is the (x, y) distance between r_i and q, and h is the high-resolution grid spacing.

Finally, we apply a bilateral filter to reduce noise while preserving sharp features. We then re-register $\{R_1, \ldots, R_N\}$ to the resulting surface approximation, and repeat the reconstruction and registration steps until convergence (three iterations in this case). Total processing time was about two minutes.

Applications

As we see in Figure 1, we can significantly improve the level of detail captured with a commercial scanner. This should be useful with valuable samples for which it is worth collecting a lot of data. The method might also inspire smaller or cheaper hardware, by passing more of the burden to software. We believe that it can be used as an alternative approach for capturing large objects at high resolution: instead of registering many high-resolution scans, we can capture many (slightly shifted) low-resolution scans covering a larger area, moving the scanner only slightly, and then merge them to produce a high-resolution surface patch.

To construct a complete model rather than a single depth image, we plan to use a modified version of Curless and Levoy's VRIP surface reconstruction [1996].

References

CURLESS, B., AND LEVOY, M. 1996. A volumetric method for building complex models from range images. *Computer Graphics 30*, Annual Conference Series, 303–312.

PARK, S. C., PARK, M. K., AND KANG, M. G. 2003. Super-resolution image reconstruction: a technical overview. *IEEE Signal Processing Magazine*, 21–36.

*e-mail: kil@cs.ucdavis.edu
†e-mail: boris@impa.br
‡e-mail: amenta@cs.ucdavis.edu