SIGGRAPH 2000 Course on
3D Photography

Shape and Appearance from Images and Range Data

Brian Curless
University of Washington

Overview

Range images vs. point clouds
Registration
Reconstruction from point clouds
Reconstruction from range images
Modeling appearance

Range images

For many structured light scanners, the range data forms a highly regular pattern known as a *range image*.

The sampling pattern is determined by the specific scanner.

Examples of sampling patterns

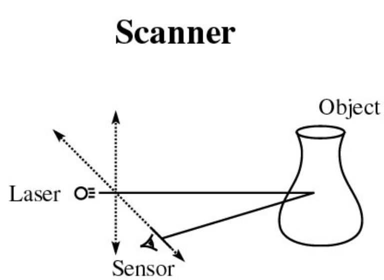

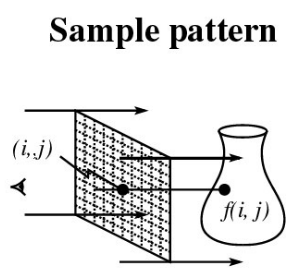

Examples of sampling patterns

Scanner

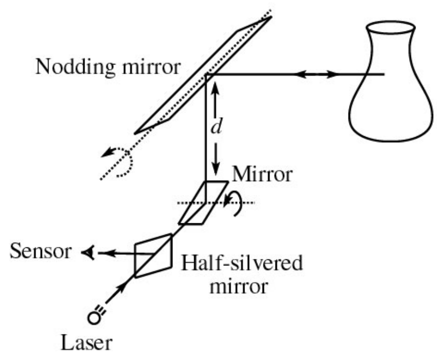

Sample pattern

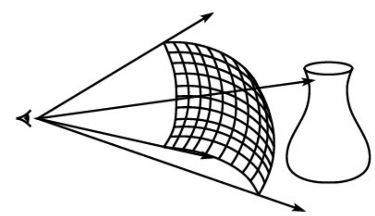

Examples of sampling patterns

Scanner

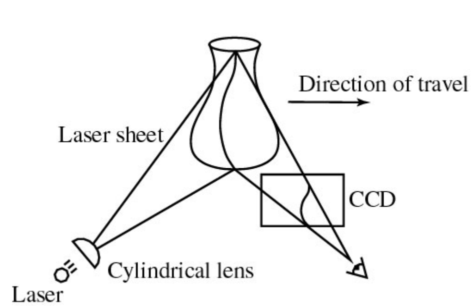

Sample pattern

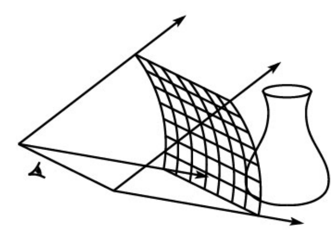

Range images and range surfaces

Given a range image, we can perform a preliminary reconstruction known as a *range surface*.

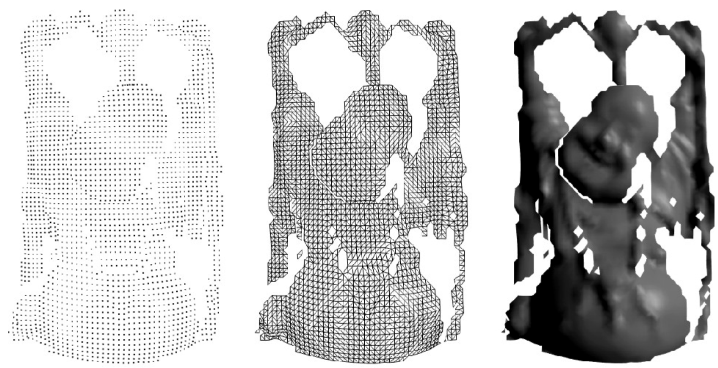

Range image Tesellation Range surface

Tessellation threshold

To avoid "prematurely aggressive" reconstruction, a tessellation threshold is employed:

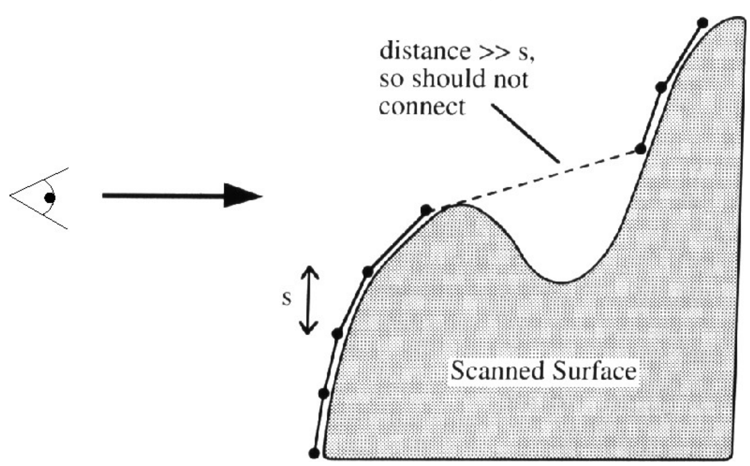

Registration

Any surface reconstruction algorithm strives to use all of the detail in the range data.

To preserve this detail, the range data must be precisely registered.

Accurate registration may require:

- ***Calibrated scanner positioning***
- ***Software optimization***
- ***Both***

Registration as optimization

Given two overlapping range scans, we wish to solve for the rigid transformation, *T*, that minimizes the distance between them.

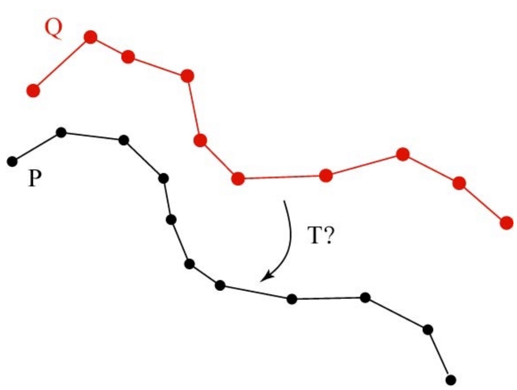

Registration as optimization

An approximation to the distance between range scans is:

$$E = \sum_{i}^{N_P} \left\| Tq_i - p_i \right\|^2$$

Where the q_i are samples from scan Q and the p_i are the *corresponding* points of scan P. These points may lay on the range surface derived from P.

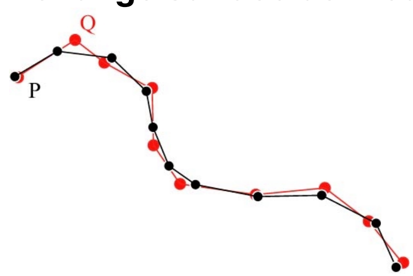

Registration as optimization

If the correspondences are known a priori, then there is a closed form solution for *T*.

However, the correspondences are not known in advance.

Registration as optimization

Iterative solutions such as [Besl92] proceed in steps:

- ***Identify nearest points***
- ***Compute the optimal T***
- ***Repeat until E is small***

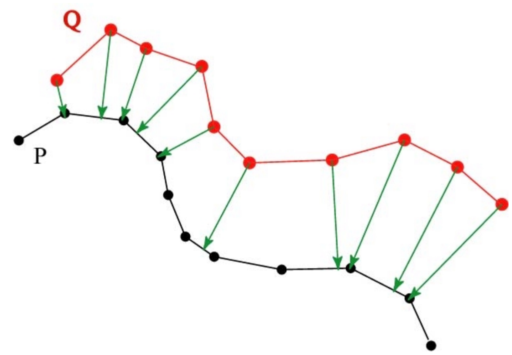

Registration as optimization

This approach is troubled by slow convergence when surfaces need to slide along each other.

Chen and Medioni [Chen92] describe a method that does not penalize sliding motions.

The Chen and Medioni method was the method of choice for pairwise alignment on the Digital Michelangelo Project.

Global registration

Pairwise alignment leads to accumulation of errors when walking across the surface of an object.

The optimal solution minimizes distances between all range scans simultaneously. This is sometimes called the *global registration* problem.

Finding efficient solution methods to the global registration problem is an active area of research.

"Non-linear registration"

Calibrating scanners can be extremely difficult.

The DMP scanner was not 100% calibrated. How to compensate?

Solution: fold non-linear scanner parameters into some of the registration procedures.

Surface reconstruction

Given a set of registered range points or range images, we want to reconstruct a 2D manifold that closely approximates the surface of the original model.

Desirable properties

Desirable properties for surface reconstruction:

- ***No restriction on topological type***
- ***Representation of range uncertainty***
- ***Utilization of all range data***
- ***Incremental and order independent updating***
- ***Time and space efficiency***
- ***Robustness***
- ***Ability to fill holes in the reconstruction***

Point clouds vs. range images

We can view the entire set of aligned range data as a point cloud or as a group of overlapping range surfaces.

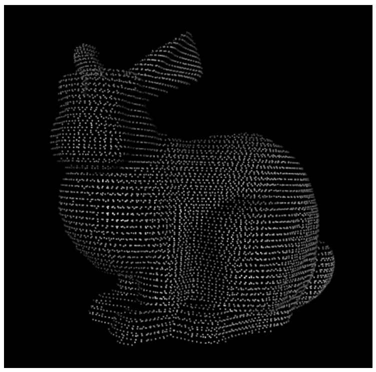

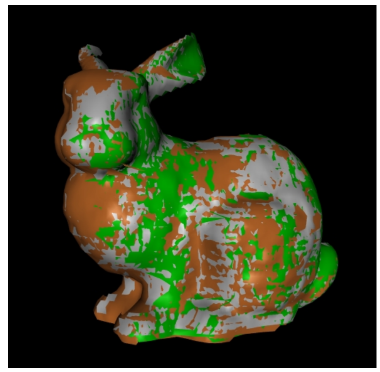

Reconstruction methods

Surface reconstruction from range data has been an active area of research for many years.

A number of methods reconstruct from unorganized points. Such methods:

- ***are general***
- ***typically do not use all available information***

Parametric vs. implicit

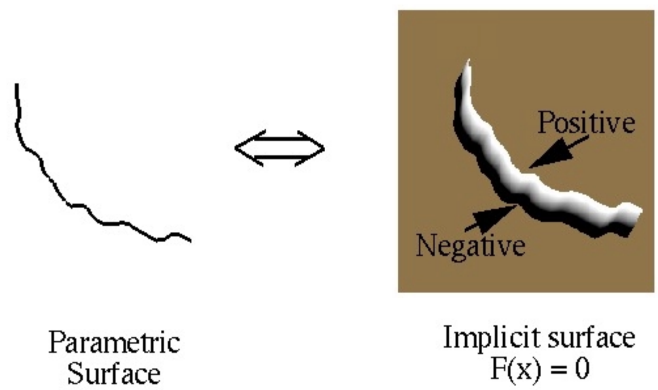

Reconstruction from unorganized points

Methods that construct triangle meshes directly:

- ***Alpha shapes [Edelsbrunner92]***
- ***Local Delaunay triangulations [Boissonat94]***
- ***Crust algorithm [Amenta98]***

Methods that construct implicit functions:

- ***Voxel-based signed distance functions [Hoppe92]***
- ***Bezier-Bernstein polynomials [Bajaj95]***

Hoppe treats his reconstruction as a topologically correct approximation to be followed by mesh optimization [Hoppe93].

Reconstruction from unorganized points

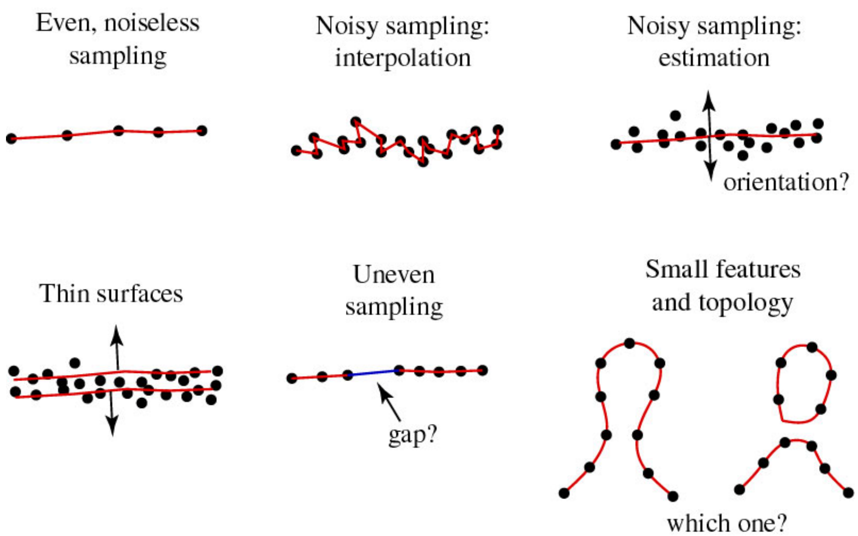

Reconstruction from range images

Methods that construct triangle meshes directly:

- ***Re-triangulation in projection plane [Soucy92]***
- ***Zippering in 3D [Turk94]***

Methods that construct implicit functions:

- ***Signed distances to nearest surface [Hilton96]***
- ***Signed distances to sensor + space carving [Curless96]***

We will focus on the two reconstruction algorithms of [Turk94] and [Curless96].

Zippering

A number of methods combine range surfaces by stitching polygon meshes together.

Zippering [Turk94] is one such method.

Overview:

- ***Tessellate range images and assign weights to vertices***
- ***Remove redundant triangles***
- ***Zipper meshes together***
- ***Extract a consensus geometry***

Weight assignment

Final surface will be weighted combination of range images.

Weights are assigned at each vertex to:

- ***Favor views with higher sampling rates***
- ***Encourage smooth blends between range images***

Weights for sampling rates

Sampling rate over the surface is highest when view direction is parallel to surface normal.

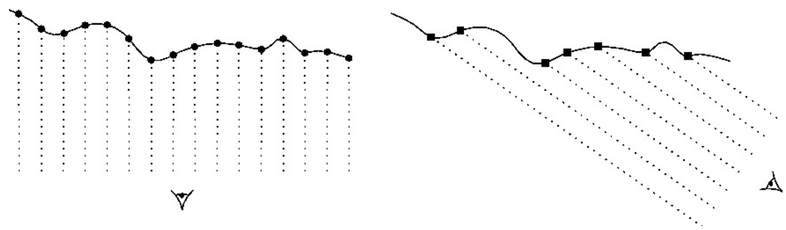

Weights for smooth blends

To assure smooth blends, weights are forced to taper in the vicinity of boundaries:

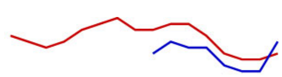

Two range surfaces

After unweighted blending

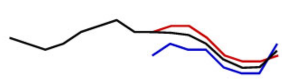

After weighted blending

Example

Range surface

Confidence rendering

Redundancy removal and zippering

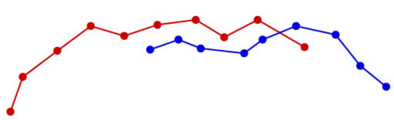

Overlapping range surfaces

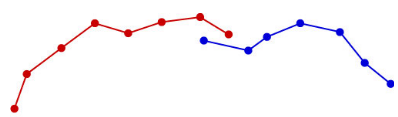

Redundant geometry removed

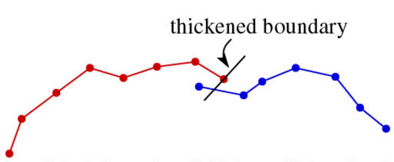

Mesh boundary "thickened" for clipping

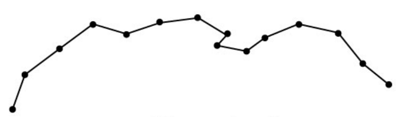

Zippered surface

Example

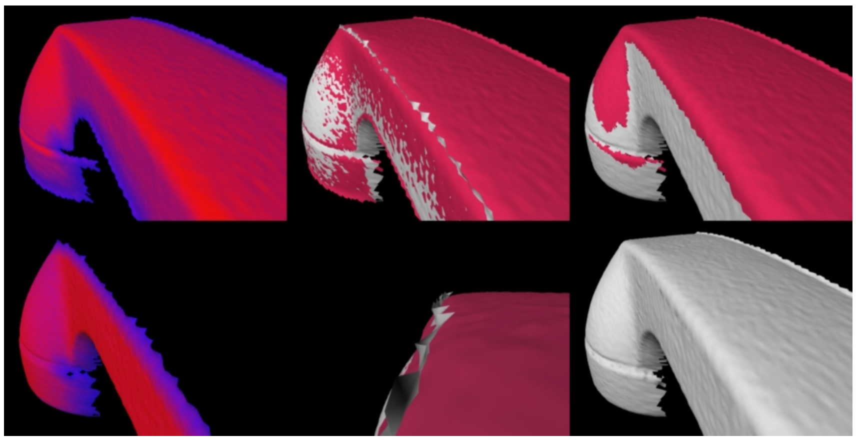

Consensus geometry

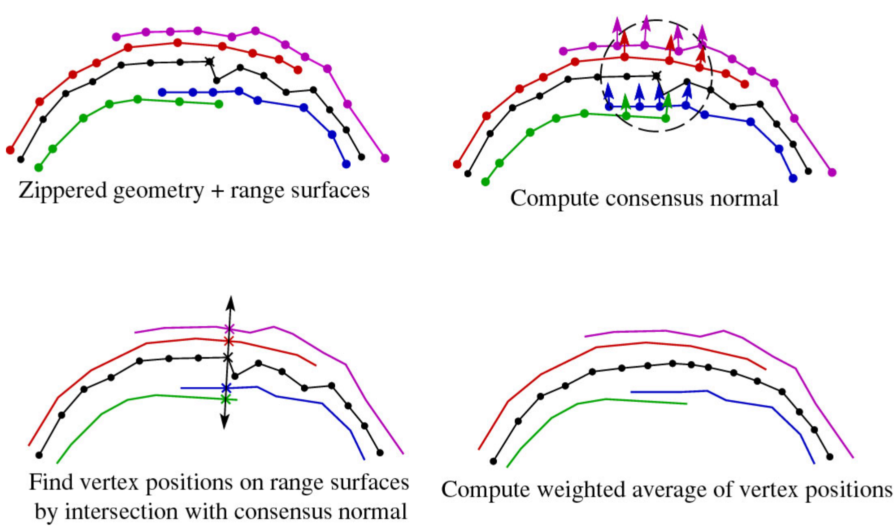

Example

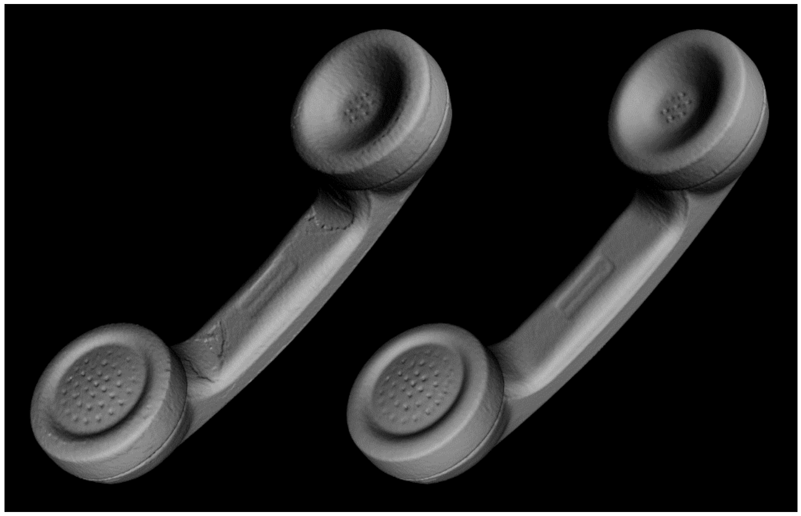

Volumetrically combining range images

Combining the meshes volumetrically can overcome some difficulties of stitching polygon meshes.

Here we describe the method of [Curless96].

Overview:

- ***Convert range images to signed distance functions***
- ***Combine signed distance functions***
- ***Carve away empty space***
- ***Extract hole-free isosurface***

Signed distance function

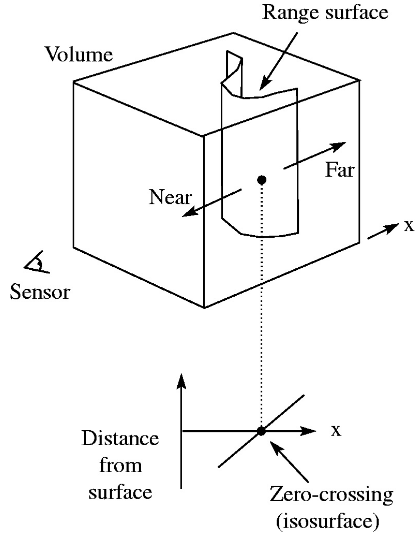

Combining signed distance functions

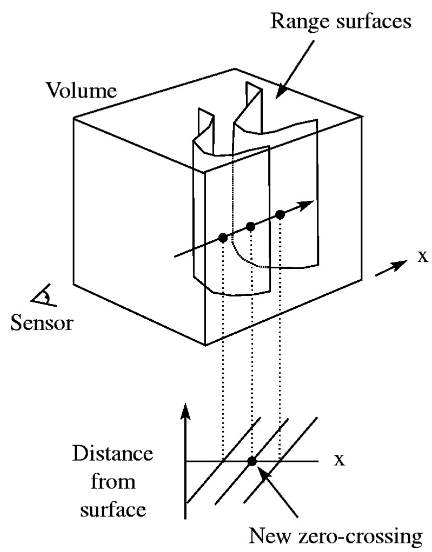

Merging surfaces in 2D

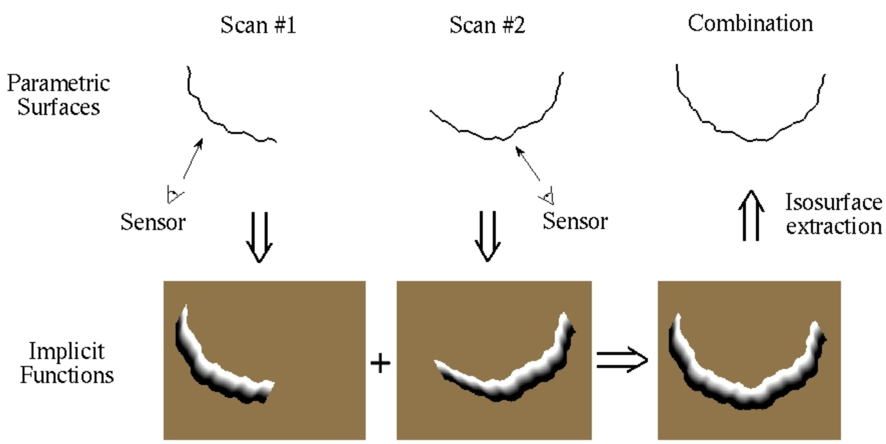

Least squares solution

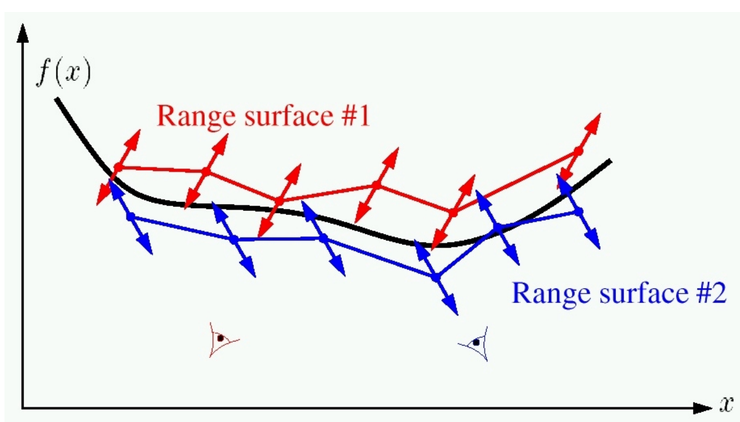

Least squares solution

$$E(f) = \sum_{i=1}^{N} \overbrace{\int \underbrace{d_i^2(x, f)}_{\text{}} dx}^{}$$

Error per point

Error per range surface

Finding the f(x) that minimizes E yields the optimal surface.

This f(x) is exactly the zero-crossing of the combined signed distance functions.

Hole filling

We have presented an algorithm that reconstructs the observed surface. Unseen portions appear as holes in the reconstruction.

A hole-free mesh is useful for:

- ***Fitting surfaces to meshes***
- ***Manufacturing models (e.g., stereolithography)***
- ***Aesthetic renderings***

Hole filling

We can fill holes in the polygonal model directly, but such methods:

- ***are hard to make robust***
- ***do not use all available information***

Space carving

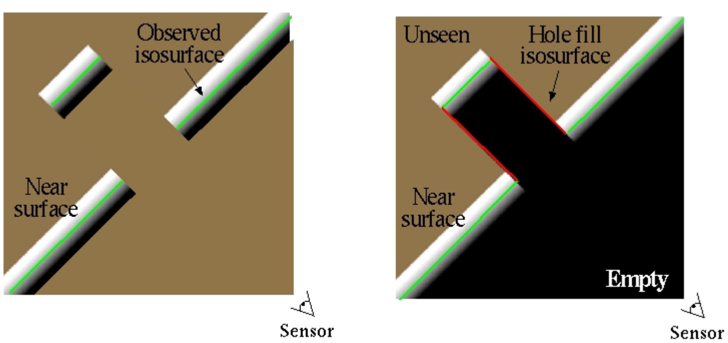

Without space carving

With space carving

Carving *without* a backdrop

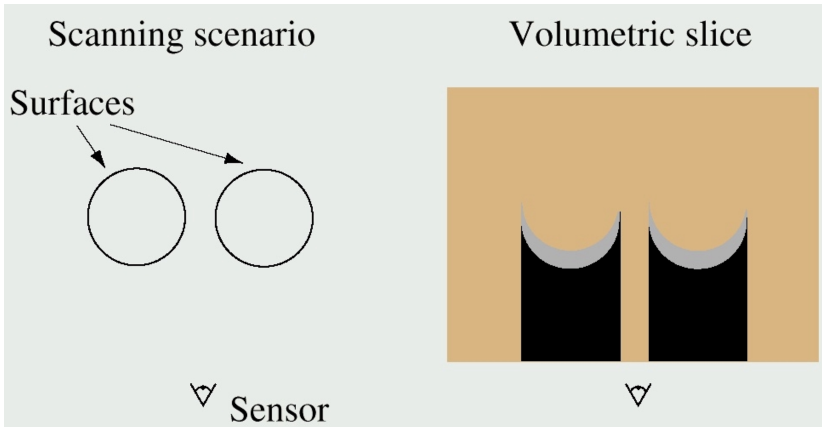

Carving *with* a backdrop

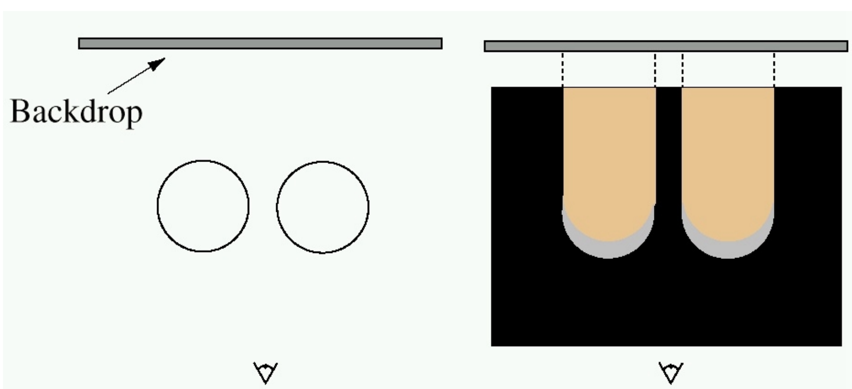

Merging 12 views of a drill bit

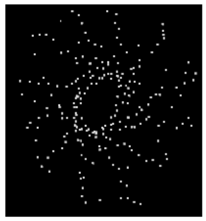

Scattered points

Range surfaces

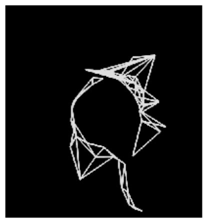

Zippered mesh

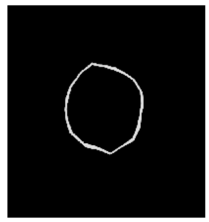

Volumetric mesh

Merging 12 views of a drill bit

Photograph of painted drill bit

Zippered mesh

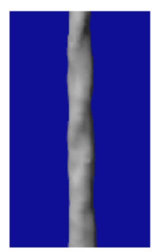

Volumetric mesh

Dragon model

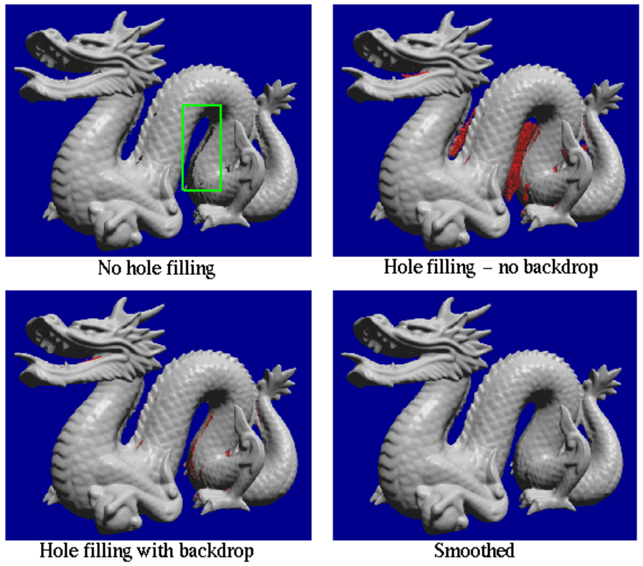

Dragon model

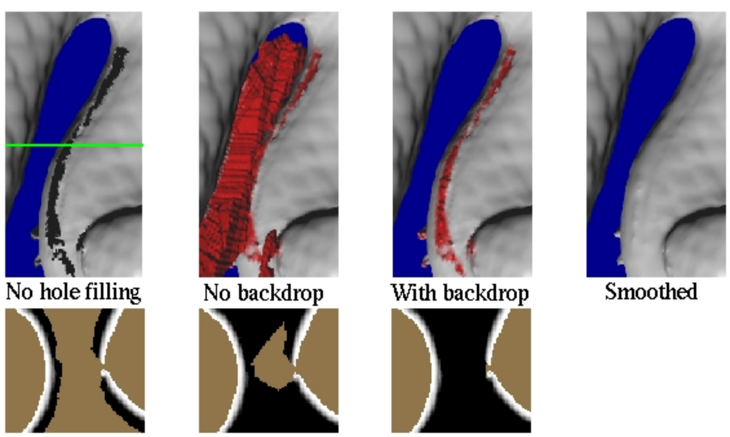

Happy Buddha

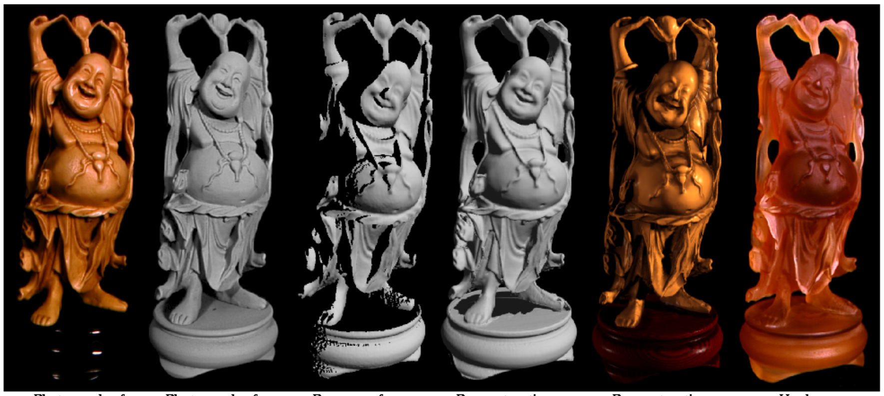

Photograph of original model	Photograph of painted original	Range surface from one scan	Reconstruction before hole–filling	Reconstruction after hole–filling	Hardcopy

Modeling appearance

When describing appearance capture, we distinguish fixed from variable lighting.

Fixed lighting yields samples of the radiance function over the surface.

This radiance function can be re-rendered using methods such as lumigraph rendering or view-dependent texture mapping.

Modeling appearance

Other methods represent, compress, and render the radiance function directly on the surface.

[Wood00] describes one such method later this week.

BRDF modeling

To re-render under new lighting conditions, we must model the BRDF.

Modeling the BRDF accurately is hard:

- ***BRDF is 4D in general.***
- ***Interreflections require solving an inverse rendering problem.***

Simplifications:

- ***Assume no interreflections***
- ***Assume a reflectance model with few parameters***

BRDF modeling

[Sato97] assume no interreflections and a Torrance-Sparrow BRDF model.

Procedure:

- ***Extract diffuse term where there are no specular highlights***
- ***Compute specular term at the specular highlights***
- ***Interpolate specular term over the surface***

BRDF modeling

[Debevec00] also develops a diffuse-specular separation technique in the context of human skin BRDF's.

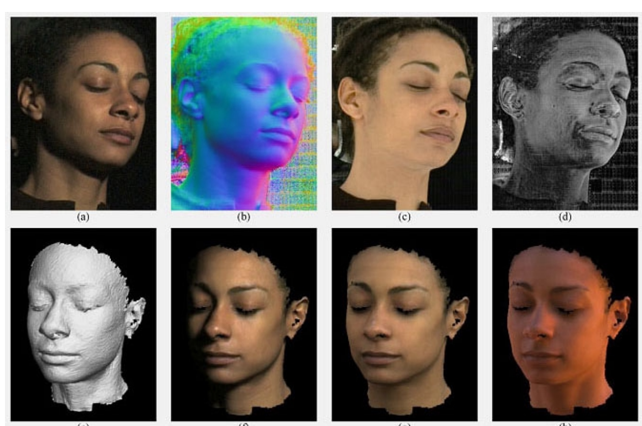

BRDF modeling

Some researchers have modeled the impact of interreflections.

[Nayar91] assumes diffuse reflectance and extracts shape and reflectance from photometric stereo.

[Yu99] has demonstrated a method that computes diffuse and specular terms given geometry, even in the presence of interreflections.

Bibliography

Amenta, N., Bern, M., and Kamvysselis, M., “A new voronoi-based surface reconstructino algorithm,” SIGGRAPH '98. Orlando, FL, USA, 19-24 July 1992. p. 415-421.

Bajaj, C.L., Bernardini, F. and Xu, G, "Automatic reconstruction of surfaces and scalar fields from 3D scans,” SIGGRAPH '95 (Los Angeles, CA, Aug. 6-11, 1995, ACM Press, pp. 109--118.

Besl, P.J. and McKay, H.D., “A method for registration of 3-D shapes,” IEEE Transactions on Pattern Analysis and Machine Intelligence, Feb. 1992, 2(4), pp. 239--256.

Boissonnat, J.-D, “Geometric Structures for Three-Dimensional Shape Representation,” ACM Transactions on Graphics, October, 1994, 3(4), pp. 266-286.

Chen, Y. and Medioni, G., “Object modeling by registration of multiple range images,” Image and Vision Computing, 10(3), April, 1992, pp. 145-155.

Curless, B. and Levoy, M., “A volumetric method for building complex models from range images.” In Proceedings of SIGGRAPH ‘96, pp. 303-312. ACM Press, August 1996.

Debevec, P., Hawkins, T., Tchou, C., Duiker, H.-P., Sarokin, W., and Sagar, M., “Acquiring the reflectance field of a human face”, SIGGRAPH ’00, pp. 145-156.

Edelsbrunner, H., Mucke, E.P. “Three-dimensional alpha shapes.” ACM Transactions on Graphics (Jan. 1994) vol.13, no.1, pp. 43-72.

Hilton, A., Stoddart, A.J., Illingworth, J., and Windeatt, T, “Reliable Surface Reconstruction from Multiple Range Images,” Fourth European Conference on Computer Vision, April 1996, vol. 1, pp. 117-126.

Bibliography

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W., "Surface reconstruction from unorganized points." SIGGRAPH '92. Chicago, IL, USA, 26-31 July 1992. p. 71-8.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W., "Mesh optimization." SIGGRAPH '93. Anaheim, CA, USA, 1-6 Aug. 1993, pp. 19-26.

Nayar, S.K., Ikeuchi, K., Kanade, T., "Recovering shape in the presence of interreflections." 1991 IEEE International Conference on Robotics and Automation, pp. 1814-19.

Sato, Y., Wheeler, M.D., Ikeuchi, K., "Object shape and reflectance modeling from observation." SIGGRAPH '97, p.379-387.

Soucy, M. and Laurendeau, D., "Multi-resolution surface modeling from multiple range views," Proceedings 1992 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, held in Champaign, IL, USA, 15-18 June 1992, pp. 348-353.

Turk, G. and Levoy, M., "Zippered polygon meshes from range images." In Proceedings of SIGGRAPH '94, pp. 311-318. ACM Press, July 1994.

Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs,"* Proc. SIGGRAPH, *Los Angeles, CA, July 1999, pp. 215-224.

Wood, D.N., Azuma, D.I., Aldinger, K., Curless, B., Duchamp, T., Salesin, D.H., Stuetzle, W., "Surface light fields for 3D photography," SIGGRAPH '00, pp. 287-296.