A Volumetric Method for Building Complex Models from Range Images

Brian Curless

Marc Levoy

Computer Graphics Laboratory
Stanford University

Introduction

Goal

Given a set of aligned, dense range images, we want to reconstruct a manifold that closely approximates the surface of the original model.

Desirable Properties

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

Previous work

From unorganized points

- Parametric (polygonal)
  - Edelsbrunner92, Boisannat94
- Implicit (volumetric)
  - Hoppe92, Bajaj95

From range surfaces

- Parametric (polygonal)
  - Turk94, Rutishauser94, Soucy95
- Implicit (volumetric)
  - Grosso88, Succi90, Hilton96

Volumetric method

- For a set of range images, $R_1, R_2, \ldots, R_N$, we construct signed distance functions $d_1(\mathbf{x}), d_2(\mathbf{x}), \ldots, d_N(\mathbf{x})$.
- We combine these functions to generate the cumulative function, $D(\mathbf{x})$.
- We extract the desired manifold as the isosurface, $D(\mathbf{x}) = 0$.

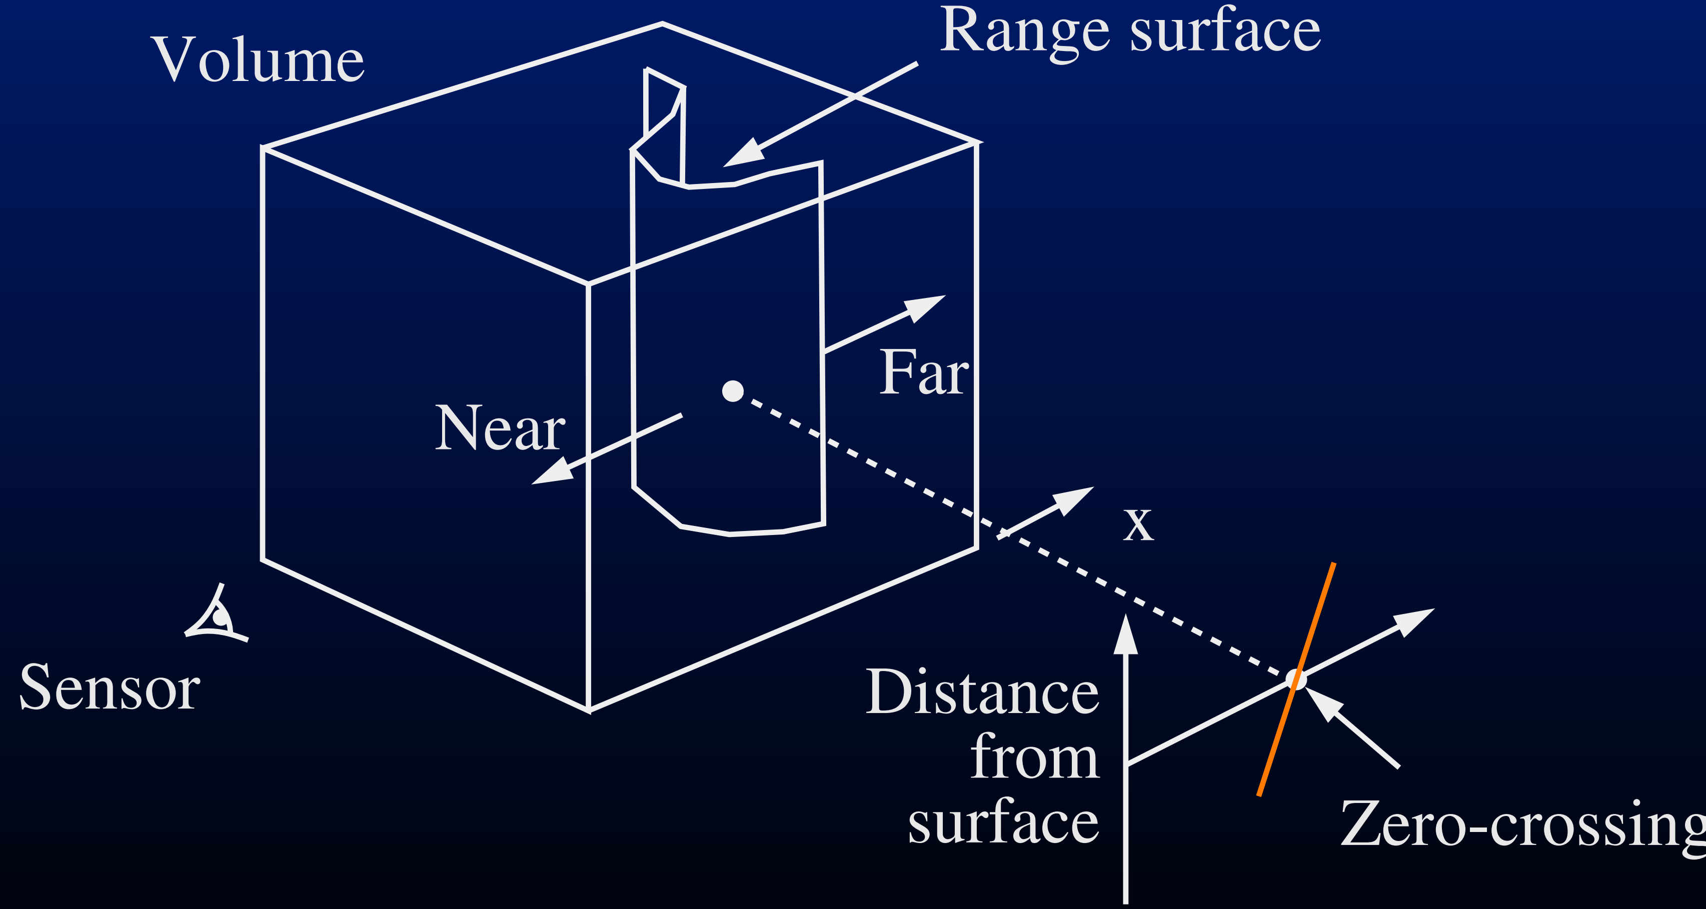
Volume
Range surface
Far
Near
x
Sensor
Distance from surface
Zero-crossing

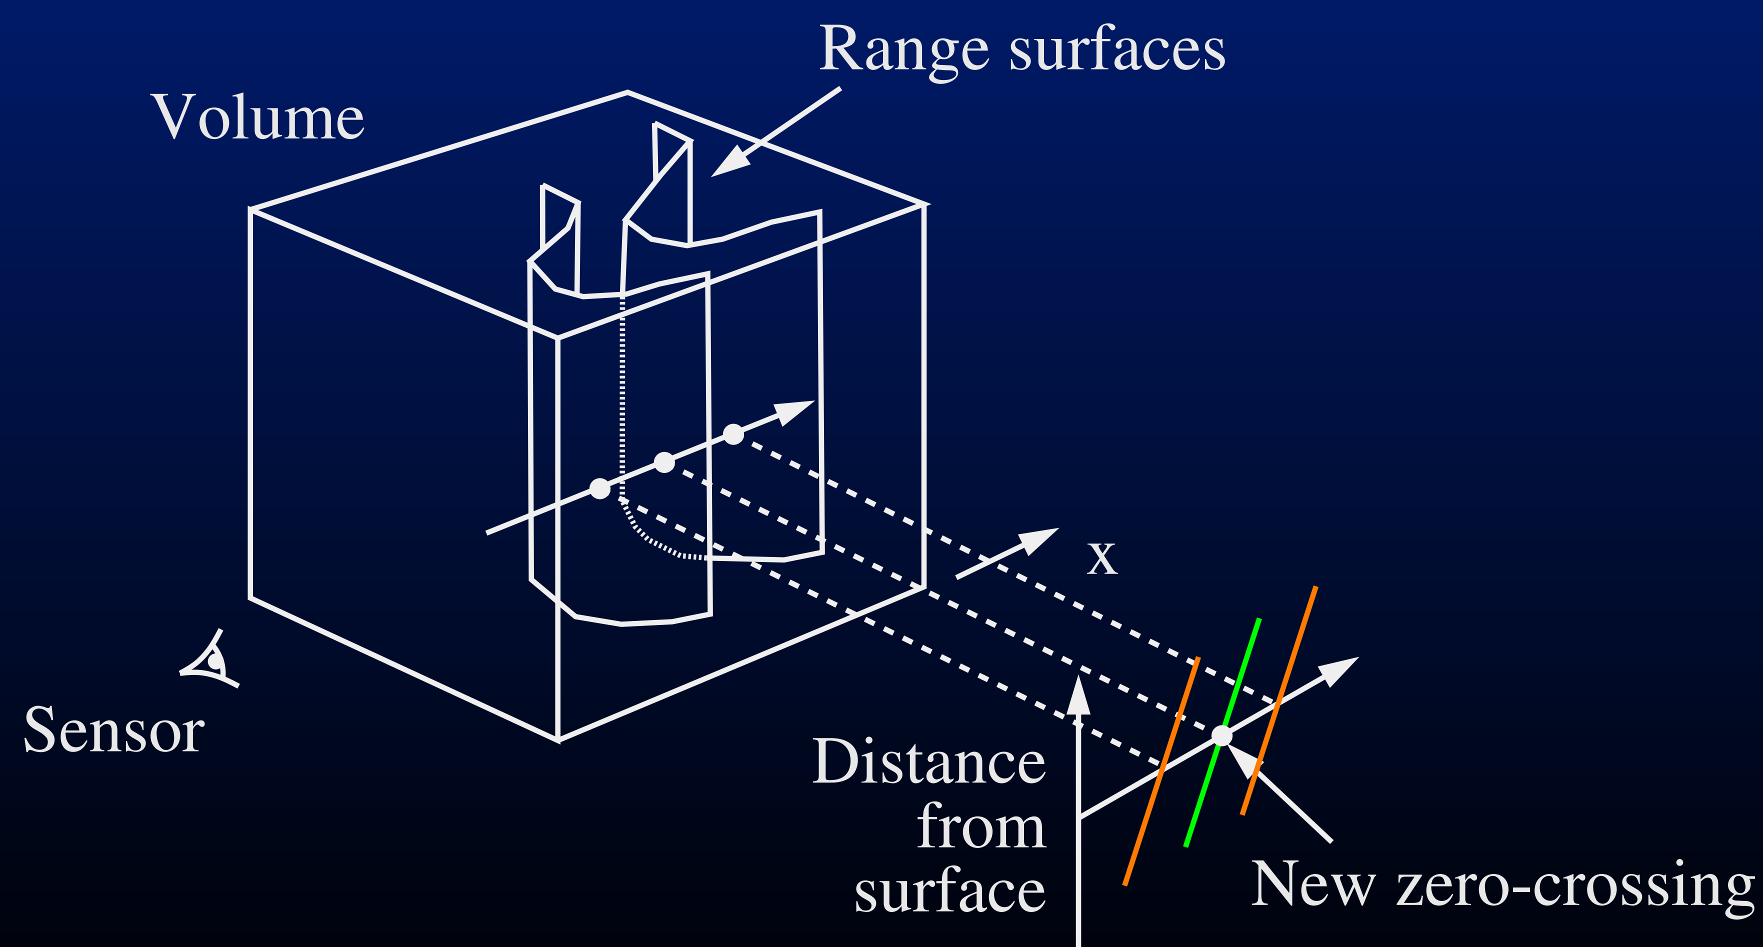
Range surfaces
Volume
x
Sensor
Distance
from
surface
New zero-crossing

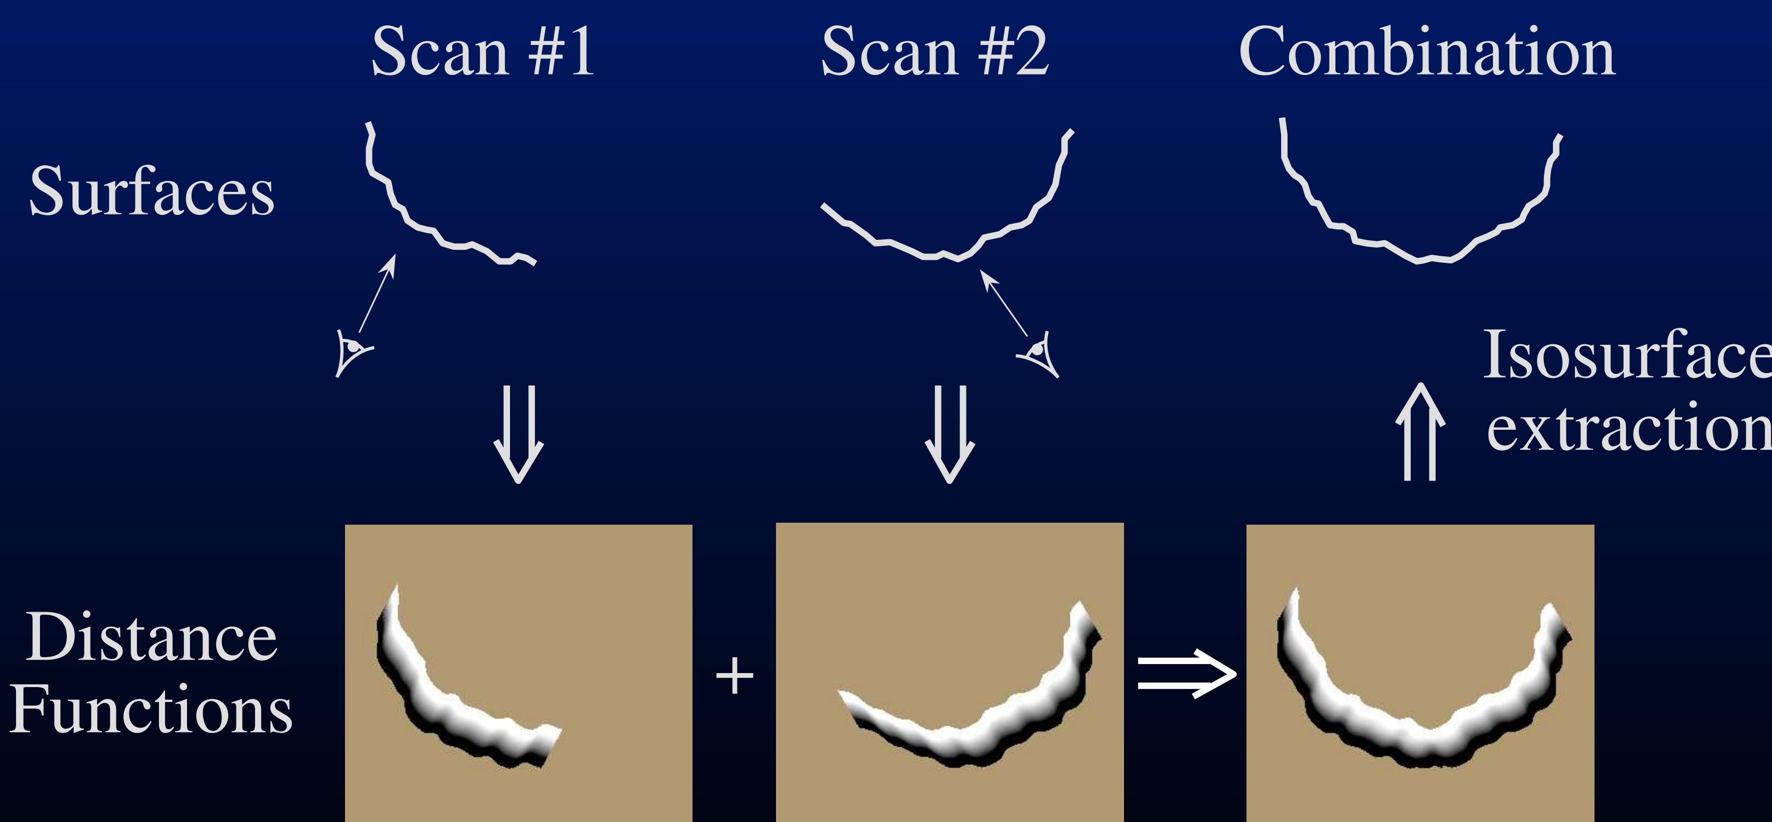

Scan #1
Scan #2
Combination
Surfaces
Isosurface extraction
Distance Functions
+

Least squares solution

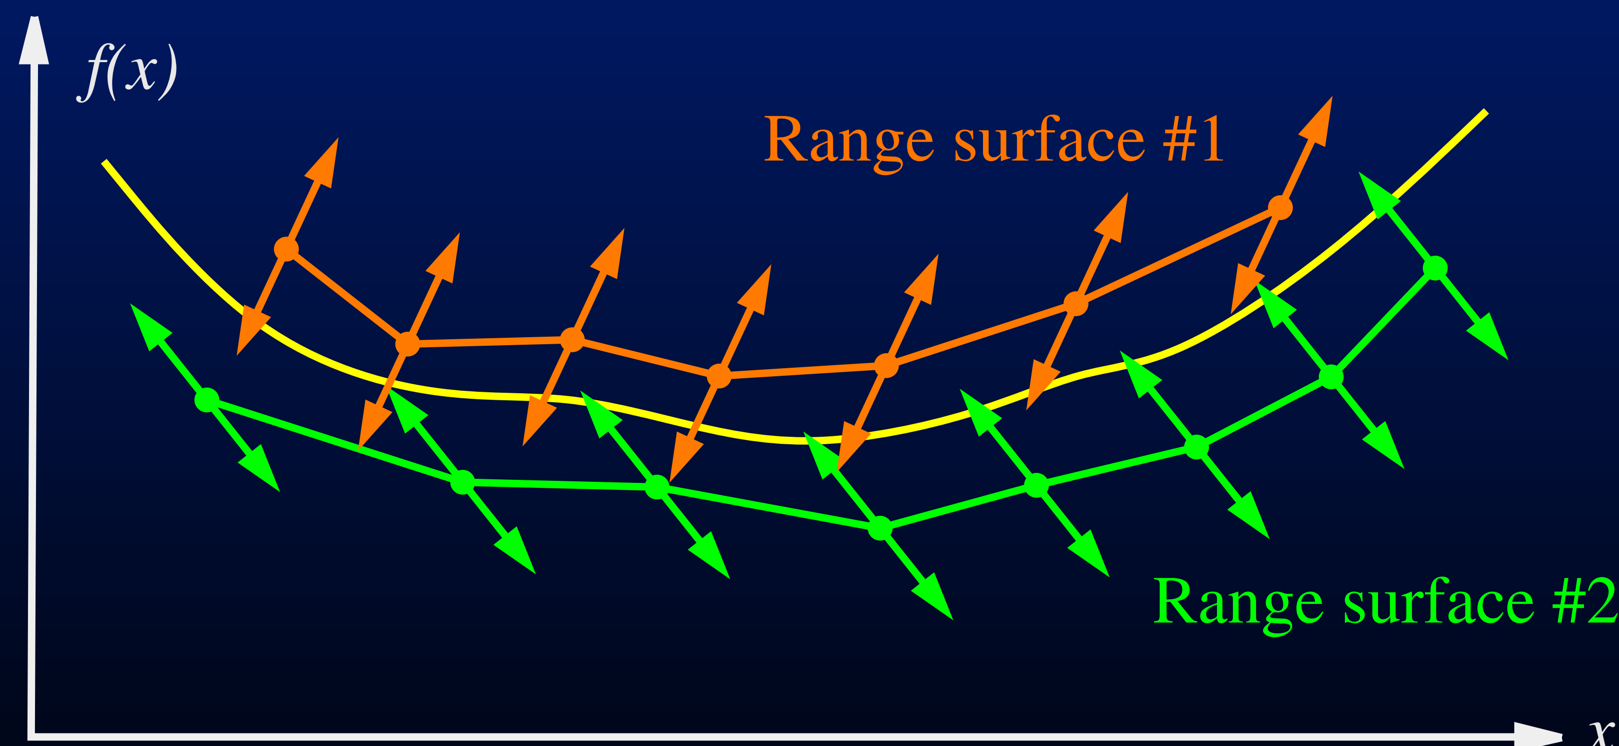

$$E(f) = \sum_{i=1}^{N} \underbrace{\int \overbrace{d_i^2(x, f)}^{\text{Error per point}} dx}_{\text{Error per range surface}}$$

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

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

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

Without space carving

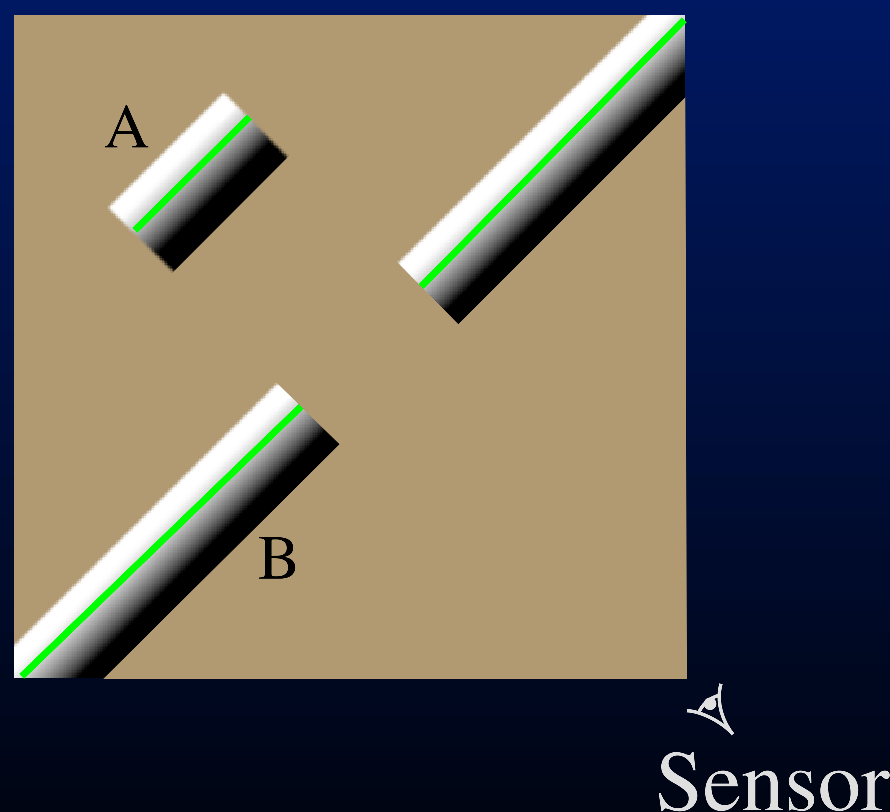

With space carving

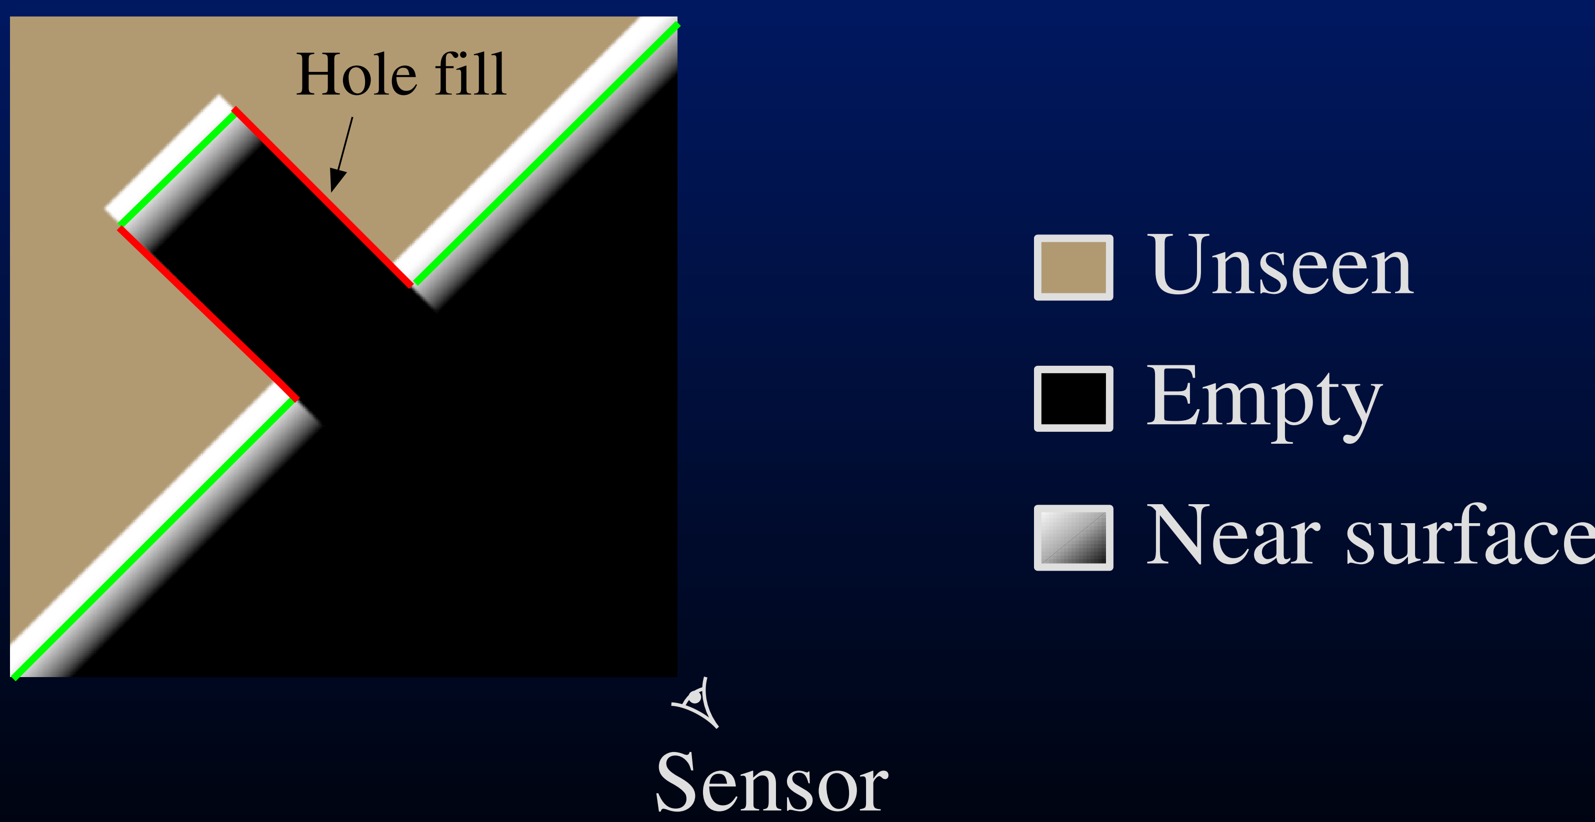

Carving *without* a backdrop

Scanning scenario

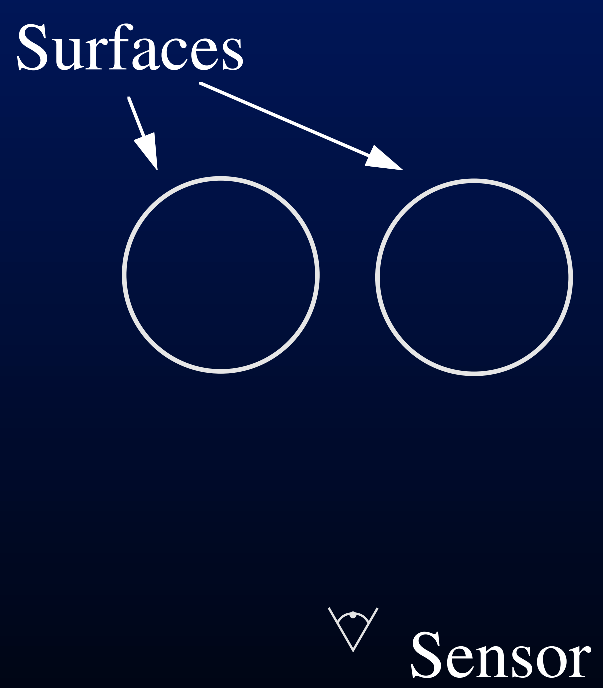

Volumetric slice

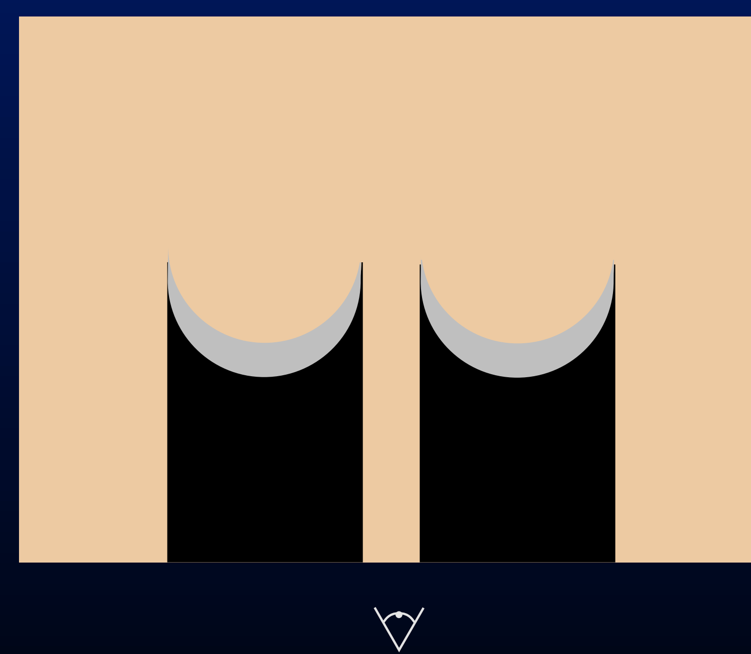

Carving *with* a backdrop

Scanning scenario

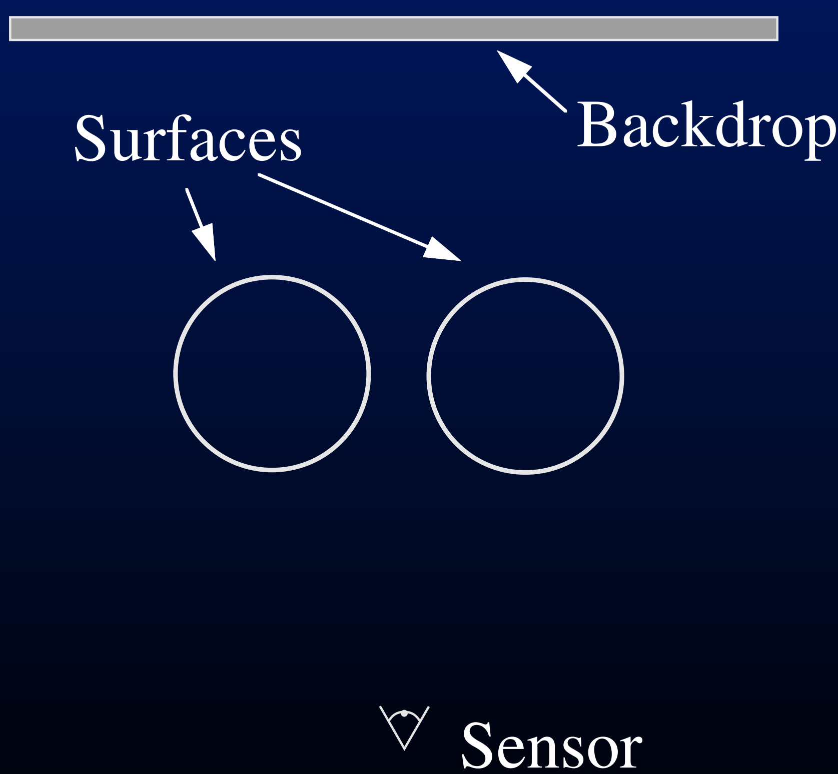

Volumetric slice

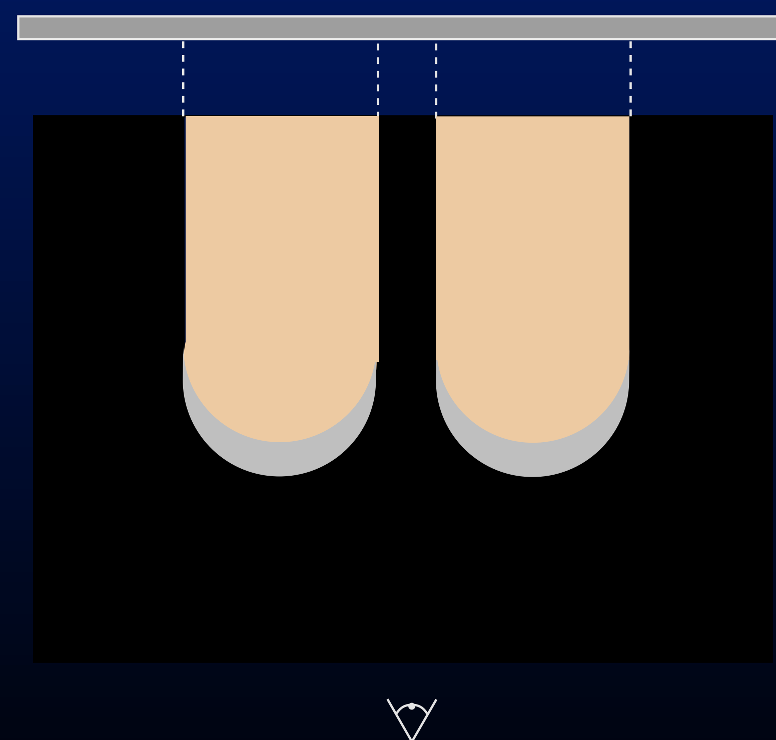

Typical data size

- 60 scans
- 10 million input vertices
- 100 million voxels

Software optimizations

- Run-length encoded data structures
- Memory coherent traversal
- Binary depth trees
- Restricted marching cubes

Results

We have tested the algorithm on several models to explore:

- Robustness (drill bit)
- Effectiveness of hole filling (dragon)
- Attainable level of detail (Happy Buddha)

Drill bit

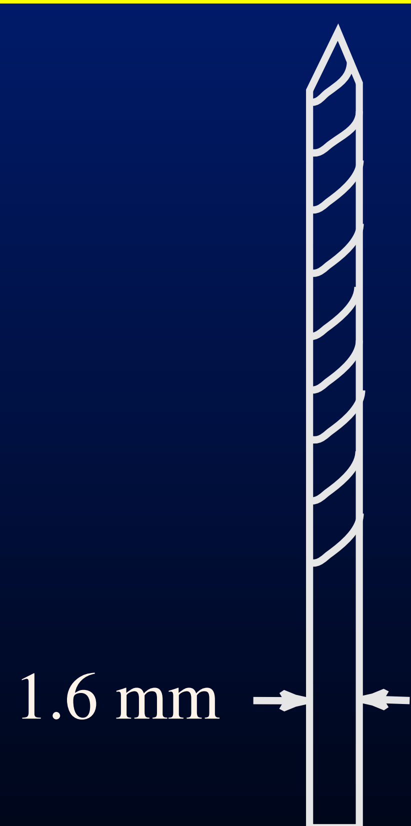

Side view of drill bit

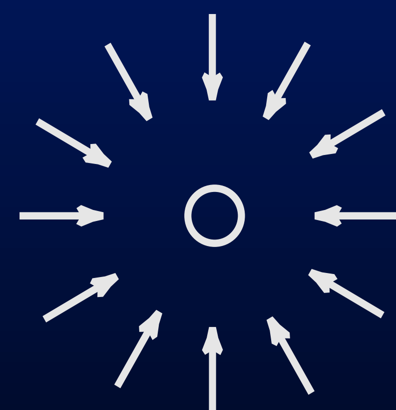

Plan view with
sensing directions

Plan view

Unorganized points

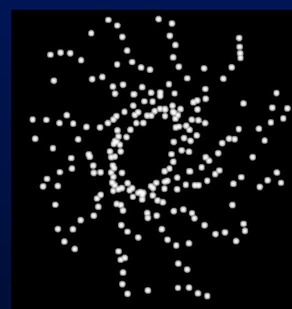

Range surfaces

Zippered mesh

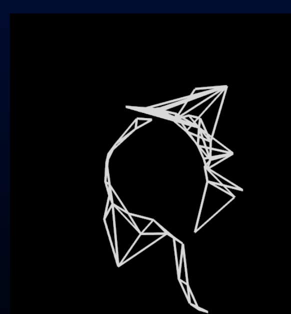

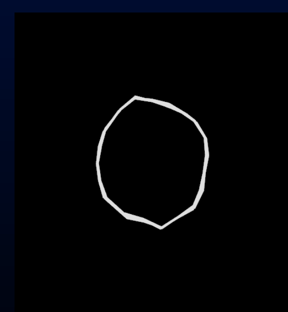

Volumetric mesh

Side view

Photograph of painted drill bit

Zippered mesh

Volumetric mesh

Dragon

1.
No hole
filling

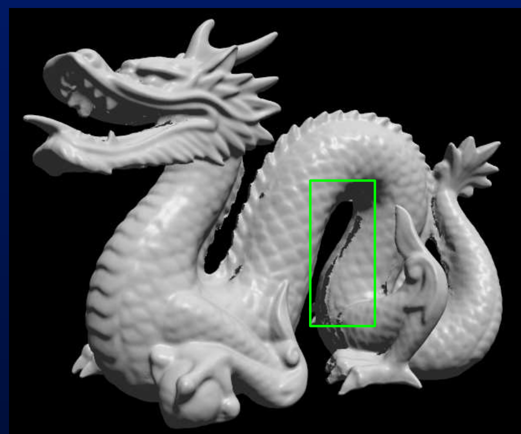

2.
Hole filling
without
backdrop

3.
Hole filling
with
backdrop

4.
Smoothed

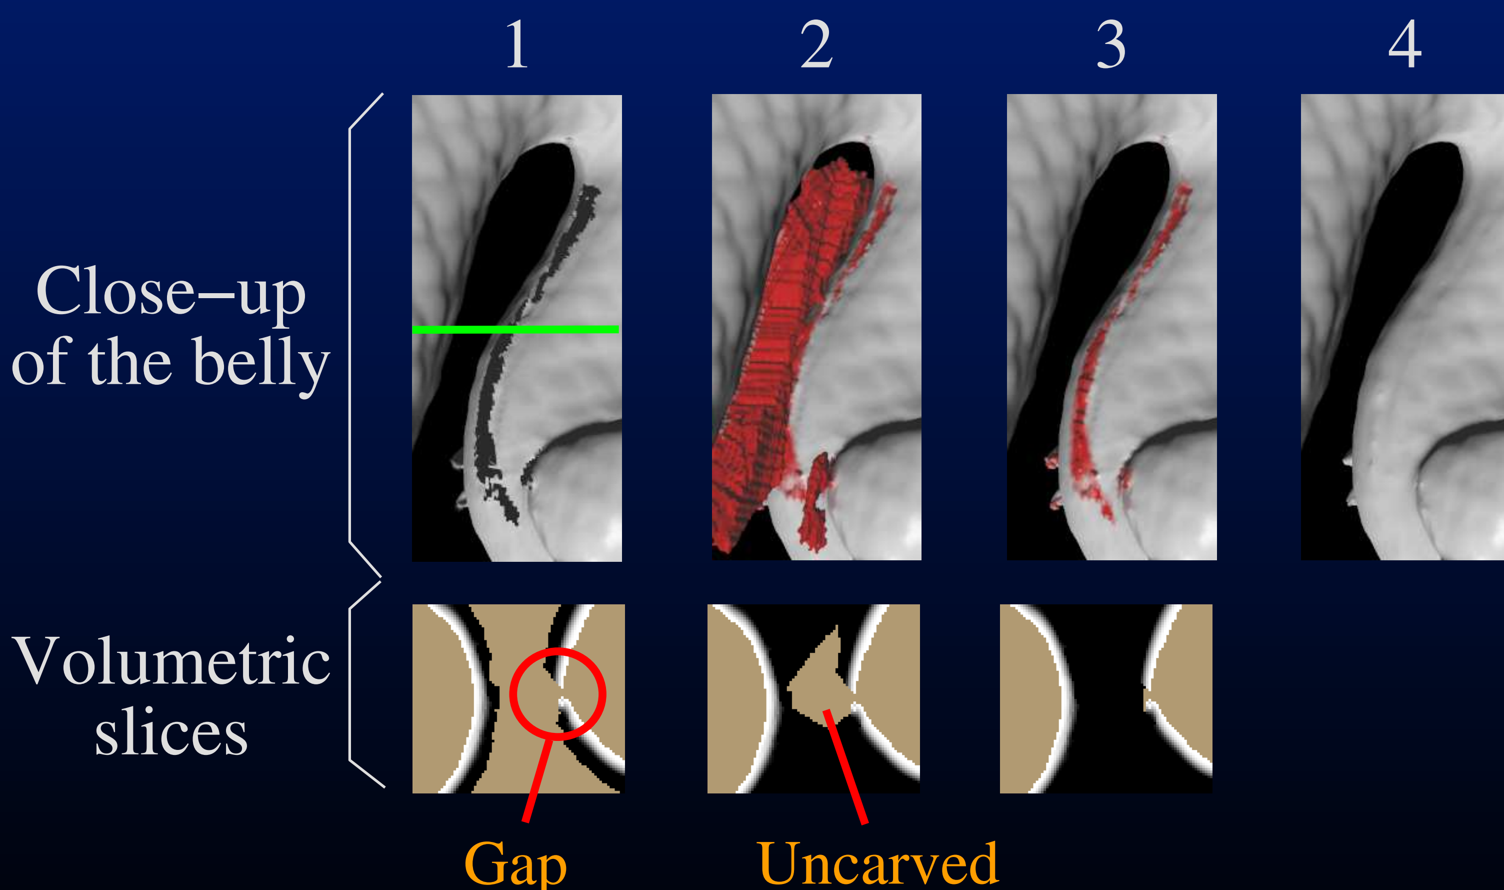

Close–up
of the belly
Volumetric
slices
1
2
3
4
Gap
Uncarved

Happy Buddha: from original...

Original model | Painted original | Range surface

...to hardcopy

Before hole filling | After hole filling | Hardcopy

Data sizes for Happy Buddha

Number of scans:	60
Input triangles:	10 million
Voxel grid:	400x1000x400
Storage:	640 MB w/o RLE
	49 MB w/ RLE
Output triangles:	2.6 million (55 MB)

Execution times for Happy Buddha

Time to scan:	1-2 hours	
Time to align:	3-4 hours	
Time to merge:	47 min.	(w/o hole fill)
	3 hr 17 min.	(w/ hole fill)
Total time:	5-10 hours	

Limitations

- Optical scanning
  - Surface points must be accessible
  - Surface reflectance affects results
- Volumetric algorithm
  - Thin surfaces and sharp corners

Future work

- Carving from video/image silhouettes
- Next best view, including backdrops
- Large-scale scenes
- Surface color acquisition

Acknowledgments

David Addleman, George Dabrowski....Cyberware scanner

Julie Dorsey, Pat Hanrahan....................Rendering tips

Homan Igehy..Triangle rasterizer

Phil Lacroute..Optimization suggestions

Bill Lorensen..Marching cubes tables

Tamara Munzner...................................Video production

Matt Pharr...Accessibility shader

Afra Zomorodian..................................Scanning script engine

Check it out...

Now available:

- Software
- Range data
- Surface reconstructions

Go to:

http://www-graphics.stanford.edu/software