Computing 3D Geometry Directly From Range Images

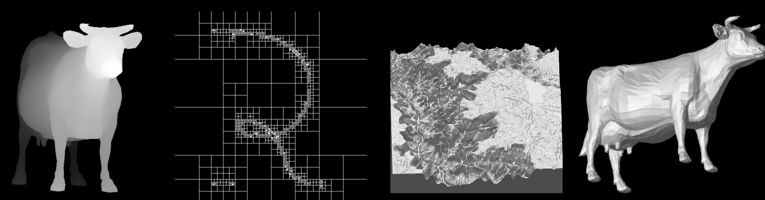

Sarah F. Frisken and Ronald N. Perry
Mitsubishi Electric Research Laboratories

Geometry from Range Data

A Classic Approach

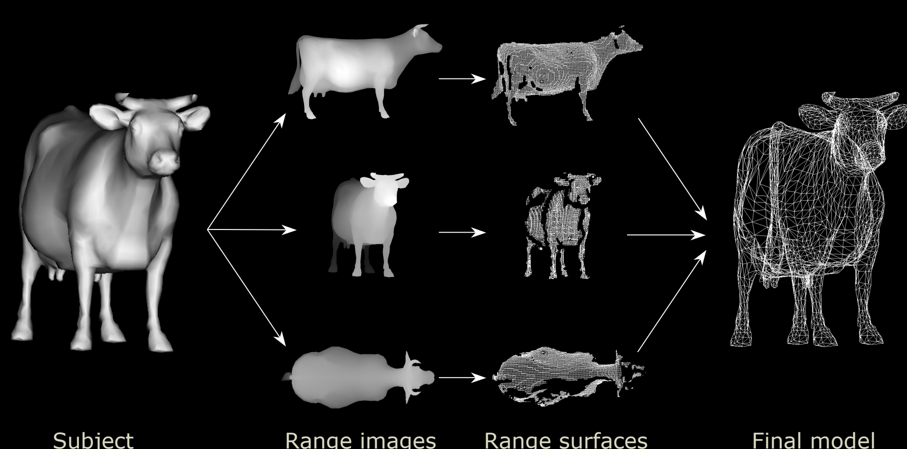

Volumetric Methods

- Create a volumetric representation of the object from the range surfaces
- Triangulate the volume data
- Advantages
  - Robust to scanner noise and image alignment errors
  - Provide good-quality, water-tight models
- Disadvantages
  - Resolution limited by volume size
  - Large memory footprint
  - Long processing times

Curless and Levoy (SIGGRAPH '96)

- Generate range surfaces from range images
- For each range surface
  - Compute signed distances to nearby volume points along the line of sight from the sensor
  - Weight the computed distance based on uncertainty
- Combine distances using a weighted average
- Triangulate the zero-valued iso-surface of the volume using Marching Cubes

Wheeler (Ph.D. thesis, CMU '96)

- Similar to Curless and Levoy except:
  - Compute *true distances* from range surfaces
  - Compute and store distances in a 3-color octree
  - Triangulate the octree using a modified Marching Cubes algorithm

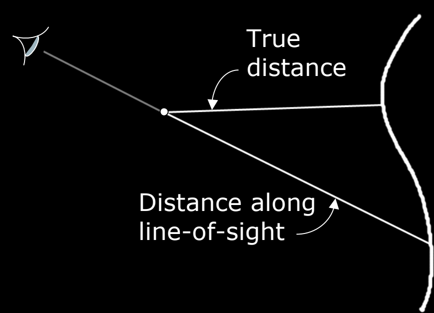

Our Approach

- Use Adaptively Sampled Distance Fields (ADFs) as the volumetric representation
  - Reduces computation and memory footprint
- Compute the ADF directly from range images
  - Avoids generating range surfaces and computing distances from range surfaces
- Add detail and edit occluded regions directly
  - ADFs provide a direct and intuitive sculpting interface
- Use a new triangulation method
  - Fast (200,000 triangles in 0.37 seconds)
  - Produces optimal triangle models
  - Produces LOD triangle models

Distance Fields

- Specify the (possibly) signed distance to a shape

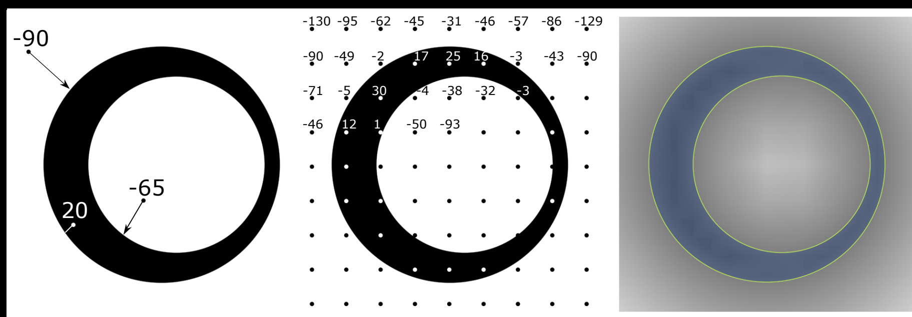

2D shape with sampled distances to its edge

Regularly sampled distance values

2D distance field

Regularly Sampled Distance Fields

- Similar to regularly sampled images, insufficient sampling of distance fields results in aliasing
- Because fine detail requires dense sampling, excessive memory is required with *regularly* sampled distance fields when *any* fine detail is present

Adaptively Sampled Distance Fields (ADFs)

- Detail-directed sampling
  - High sampling rates only where needed
- Spatial data structure (e.g., an octree)
  - Fast localization for efficient processing
- Reconstruction method (e.g., trilinear interpolation)
  - For reconstructing the distance field and its gradient from the sampled distance values

Advantages of ADFs

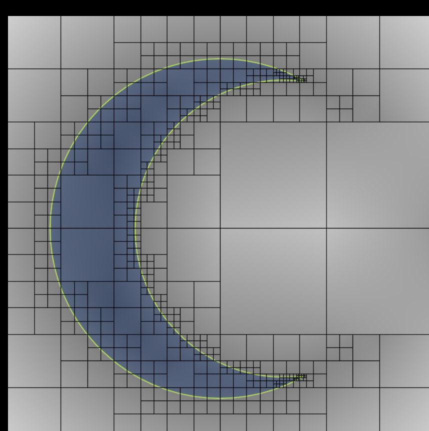

ADFs consolidate the data needed to represent complex objects

ADFs provide

- Spatial hierarchy
- Distance field
- Object surface
- Object interior
- Object exterior
- Surface normal (gradient at surface)
- Direction to closest surface point (gradient off surface)

Comparison of 3-color Quadtrees and ADFs

- Fewer distance computations
- Smaller memory footprint

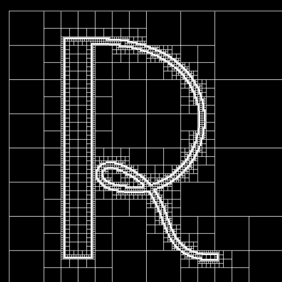

23,573 cells (3-color)

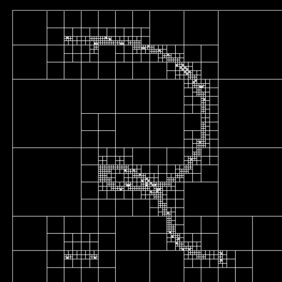

1713 cells (ADF)

Computing ADFs Directly from Range Images

- Range images measure distances along the line-of-sight rather than true distances
- At each ADF sample point
  - Compute the projected distance to the object surface from the line-of-sight distance for each range image
  - Correct each projected distance by dividing by the local gradient magnitude of the projected distance field to approximate the true distance
  - Choose the true distance with the highest confidence
- The local gradient magnitude is constant in the direction perpendicular to the range image
  - Can be pre-computed and stored as a 2D image

Editing Occluded Regions and Adding Detail to Scanned Models

- ADFs are a volumetric representation
  - Provide an intuitive interface for direct sculpting of 3D models
- Kizamu (Perry and Frisken, SIGGRAPH 2001)
  - System for sculpting digital characters
  - Can sculpt high resolution ADFs (equivalent to 2048^3 volumes) at interactive rates
  - Reasonable memory footprint
  - Produces LOD triangulations of the sculpted models

Triangulation Method

- Seed
  - Each boundary leaf cell of the ADF is assigned a vertex that is initially placed at the cell's center
- Join
  - Vertices of neighboring cells are joined to form triangles
- Relax
  - Vertices are moved to the surface using the distance field
- Improve
  - Vertices are moved over the surface towards their average neighbors' position to improve triangle quality

Creating LOD Triangle Models

- Adapt triangulation to generate LOD models
  - Traverse octree from root to leaf cells
  - Seed vertices in (possibly) non-leaf boundary cells that satisfy a minimum error criterion
  - Ignore cells below these in the hierarchy

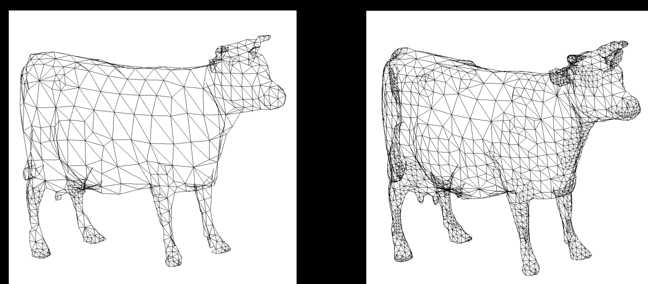

Results

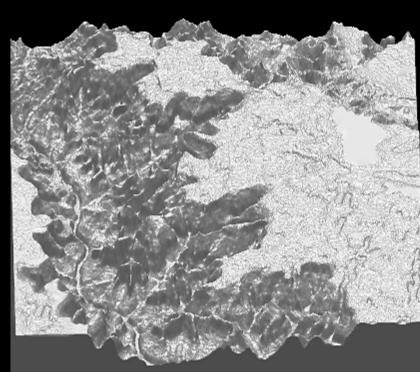

ADF generated from an 800x800 elevation image of the Grand Canyon

Results

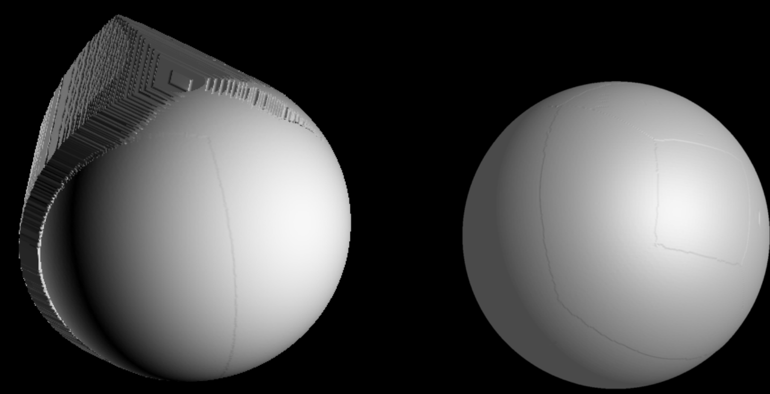

Sphere generated from 2 range images (synthetic data)

Sphere generated from 14 range images (synthetic data)

Results

Two views of a cow model generated from 14 range images (synthetic data)

Summary

- Use of distance fields provides more robust methods and water-tight surfaces
- ADFs result in significant savings in memory and distance computations
- Distances are computed directly from range images rather than from range surfaces
- Resultant models can be directly sculpted to add detail and to edit occluded regions
- Fast new triangulation method produces optimal triangle meshes from the ADF

Future Work

- Address noise in scanned data
  - Incorporate probabilistic methods from prior art for combining multiple scans
  - Extend probabilistic methods to exploit the distance field
- Align scans using the distance field (replacing point-based alignment methods such as Iterative Closest Point)

The End