

Computing 3D Geometry Directly from Range Images

Sarah F. Frisken, Ronald N. Perry

TR2001-10 December 2001

Abstract

Several techniques have been developed in research and industry for computing 3D geometry from sets of aligned range images. Recent work has shown that volumetric methods are robust to scanner noise and alignment uncertainty and provide good quality, water-tight models. However, these methods suffer from limited resolution, large memory requirements and long processing times, and they produce excessively large triangle models. In this report, we propose a new volumetric method for computing geometry from range data that: 1) computes distances directly from range images rather than from range surfaces, 2) generates an Adaptively Sampled Distance Field (ADF) rather than a distance volume or a 3-color octree, resulting in a significant savings in memory and distance computations, 3) provides an intuitive interface for manually correcting the generated ADF, and 4) generates optimal triangle models (with fewer triangles in flat regions and more triangles where needed to represent surface detail) from the generated ADF octree using a fast new triangulation method.

SIGGRAPH 2001 Conference Abstracts and Applications

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Computing 3D Geometry Directly from Range Images

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

Introduction

Several techniques have been developed in research and industry for computing 3D geometry from sets of aligned range images [1]. Recent work has shown that volumetric methods are robust to scanner noise and alignment uncertainty and provide good quality, water-tight models [2,3,4]. However, these methods suffer from limited resolution, large memory requirements and long processing times, and they produce excessively large triangle models.

The methods of [2,3,4] construct range surfaces for each aligned range image and fill a (fixed resolution) volumetric representation with signed distances from the range surfaces. The methods use various approaches to reduce the time required to fill and access this volume data, including run length encoding of the distance values, binary encoding of regions outside a bounded region of the surface, and a 3-color octree representation of the volume. The distance values from multiple scans are combined probabilistically using order-independent or incremental updating. Finally, these methods build a triangle model of the iso-surface of the distance volume using Marching Cubes.

In this sketch, we propose a new volumetric method for computing geometry from range data that: 1) computes distances directly from range images rather than from range surfaces, 2) generates an Adaptively Sampled Distance Field (ADF) rather than a distance volume or a 3-color octree, resulting in a significant savings in memory and distance computations, 3) provides an intuitive interface for manually correcting the generated ADF, and 4) generates optimal triangle models (with fewer triangles in flat regions and more triangles where needed to represent surface detail) from the generated ADF octree using a fast new triangulation method.

Corrected, Projected Distance Images

Constructing 3D range surfaces and computing distances from these surfaces contribute significantly to the computational requirements of [2,3,4]. If, instead, the distance field could be generated directly from 2D range images, model generation times could be reduced. However, range images do not provide true distance data. In the simplest case, a range image records the perpendicular projected distance from the object surface to the image plane. The projected distance field is the same as the true distance field in two circumstances: 1) *throughout the field* for a planar surface parallel to the image plane, and 2) *at the surface* (where both distances are zero) for any surface. Other than case 1), the projected distance field differs from the true distance field for points off the surface, resulting in artifacts when combining projected distance fields from different viewpoints ([2] suffers from this problem).

It can be shown mathematically for a planar surface that the difference between the true distance and the projected distance at a location \mathbf{x} is inversely proportional to the magnitude of the distance field gradient at \mathbf{x} when the gradient is computed using central differences. Here we propose to *correct* the 3D projected distance field by dividing sampled distances by the local gradient magnitude. This results in a better approximation of the true distance field near the surface, yielding better results when combining projected distance fields (see Figures 1 and 2). Computing the local 3D gradient to make this correction could be prohibitive (it requires 6 additional distance computations). Instead, we derive the 3D gradient from a 2D gradient image generated once during preprocessing, resulting in significantly faster generation.

Since the range images of many scanning systems are not simple projected distances, they must be converted to this form. However, we have found that conversion is possible from many formats (e.g. laser striping requires a 1D scan-line conversion). While this conversion results in some loss of information from the scan, the benefit is faster performance, allowing interactive updating during data acquisition.

Adaptively Sampled Distance Fields

We recently proposed adaptively sampled distance fields (ADFs) as a new representation for shape [5]. ADFs adaptively sample the signed distance field of an object and store the sample values in a spatial hierarchy (e.g., an octree) for fast processing. ADFs are memory efficient and detail directed, so that distance values are computed from the range images only where needed (i.e. mostly near highly detailed regions of the surface). [5] found that even in 2D, ADFs require 20x fewer distance computations than a comparable 3 color octree representation (used in [4]). Finally, ADFs can be interactively edited via a sculpting interface [6] so that holes and other surface anomalies from occlusions and sensor noise can be easily corrected.

The ADF is generated from sequential or order-independent range images using the tiled generator presented in [6]. Currently, distance values from the range images are combined as though carving the shape from a solid cube of material

using a Boolean differencing operator; we have also begun experimenting with adding the probabilistic combining functions of [2,3,4] for robustness to sensor noise.

Fast, Optimal Triangulation

One of the advantages of merging range images in a volumetric distance field is that the field's iso-surface yields a water-tight model. However, [2,3,4] rely on the Marching Cubes algorithm to triangulate the iso-surface which requires that all surface cells be at the same resolution, thus producing an excessive numbers of triangles. Instead, we take advantage of a new algorithm for triangulating the ADF octree [6], automatically generating fewer triangles in flat regions of the surface and more triangles where the surface has high detail. This method tends to generate an order of magnitude fewer triangles than Marching Cubes for range data, is very fast (generating models with more than 200,000 triangles in 0.37 seconds), and can produce level-of-detail models ideal for applications such as games and physical simulations.

References

- [1] B. Curless, "From range scans to 3D models", *Comp. Graphics*, Vol.33, No.4, 1999.
- [2] B. Curless and M. Levoy, "A volumetric method for building complex models from range images", *Proc. SIGGRAPH'96*, pp. 303-312, 1996.
- [3] A. Hilton, A. Toddart, J. Illingsworth, and T. Winder, "Reliable surface reconstruction from multiple range images", 4th Euro. Conf. Comp. Vis., pp. 117-126, 1996.
- [4] M. Wheeler, "Automatic modeling and localization for object recognition", Ph.D. Thesis, Carnegie Mellon University, 1996.
- [5] S. Frisken, R. Perry, A. Rockwood, and T. Jones, "Adaptively sampled distance fields: a general representation of shape for computer graphics", *Proc. SIGGRAPH'00*, pp. 249-254, 2000.
- [6] R. Perry and S. Frisken, "Kizamu: A system for sculpting digital characters", *Proc. SIGGRAPH'01*, 2001.




Figure 1. Left: *True* gradient vectors (red) of the distance field of a circle overlay gradient vectors (white) *computed* from the projected distance field measured from the left side of the image. Right: Gradient vectors computed from the *corrected* projected distance field (white) more closely match the true gradient vectors (red), especially near the circle's silhouette.




Figure 2. Left: An ADF generated from two synthetic, projected range images of a sphere (at 0 and 90 deg.) has artifacts on the seam between the two views. Right: The artifacts are greatly reduced when the *corrected* projected range images are used. Generation time: 6 seconds.


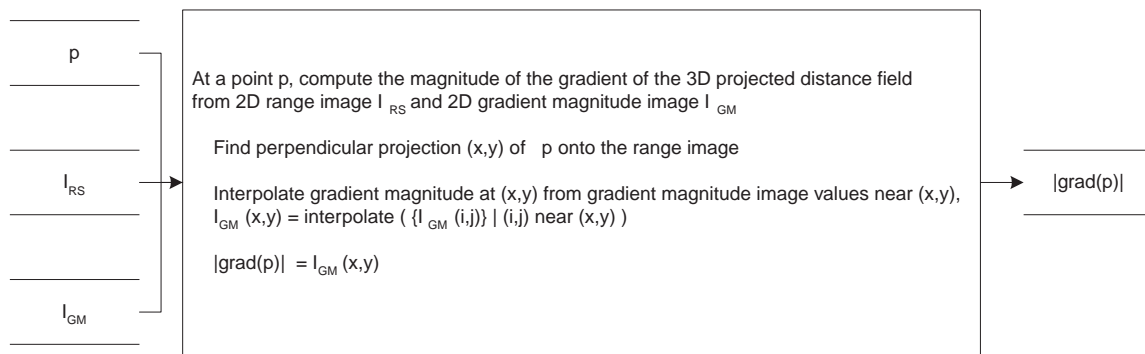
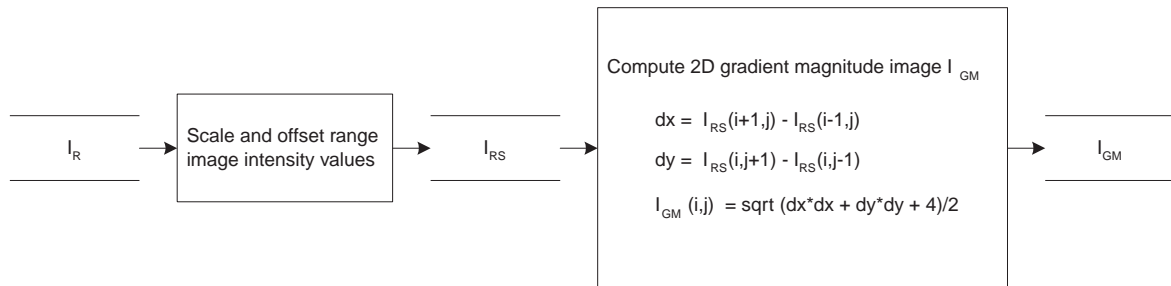


Figure 3. An ADF generated from an 800x800 elevation image of the Grand Canyon (Data courtesy of USGS). Generation time (from range image to rendered model): 15 seconds.



A data flow diagram for computing the magnitude of a 3D gradient from a 2D gradient image