

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE September 19, 2003	3. REPORT TYPE AND DATES COVERED Final, Oct 2000 - Dec 2003; September 2003	
4. TITLE AND SUBTITLE A Statistical Approach to 3D Terrain Reconstruction		5. FUNDING NUMBERS N00014-01-1-0033 A0001	
6. AUTHORS Ross T. Whitaker			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Utah School of Computing 50 S Central Campus Dr Rm 3190 Salt Lake City, UT 84112-9205		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Regional Office Seattle 1107 NE 45 th Street Suite 350 Seattle, WA 98105-4631		10. SPONSORING / MONITORING AGENCY REPORT NUMBER ONR	
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for Public Release; distribution is Unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (<i>Maximum 200 words</i>) This report summarizes the work and accomplishments of the University of Utah on the ONR grant #N00014-01-1-0033 entitled "A Statistical Approach To 3D Terrain Reconstruction". The report, as agreed with the Program Officer, Dr. L. Rosenblum, is the same as the presentation of the final meeting for this project, and consists of a series of slides with notes that explain the meaning of each slide. The report describes the motivation and overall goals of the project, presents results from this work, lists the publications that have resulted from this grant, and describes the mechanisms by which technology has been disseminated. This project addresses the question of surface reconstruction, particularly with the application of visualizing and analyzing terrain data. Results include fundamental developments in how to formulate the problem of surface fitting and smoothing and engineering developments, including fast algorithms on specialized hardware.			
14. SUBJECT TERMS Surface reconstruction, surface estimation, surface processing, laser range image, LADAR, level sets, tomography, volume graphics, partial differential equations, deformable models, surface segmentation, graphics processing cards, streaming architectures		15. NUMBER OF PAGES 42	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT none	18. SECURITY CLASSIFICATION OF THIS PAGE none	19. SECURITY CLASSIFICATION OF ABSTRACT none	20. LIMITATION OF ABSTRACT none

NSN 7540-01-280-5500

BEST AVAILABLE COPY

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-1
298-102

20040615 117

Surface Estimation and Analysis

Ross Whitaker
SCI Institute, School of Computing
University of Utah
ONR Final Report
September 19, 2003

Whitaker, ONR Sept. 19, 2003

Motivation

- Surface Measurement Technologies
 - LADAR, structured light, vision
 - Tomography
 - Sonar/ultrasound
 - Radar
- New computational capabilities
- Applications
 - Visualization – E.g command and Control
 - Recognition – Surveillance
 - Analysis

Whitaker, ONR Sept. 19, 2003

This work addresses a changing technology landscape. There are lots of new and cheaper technologies for measuring surface shape, and some older technologies (such as MRI and tomography) can be treated as surface measurement mechanisms under certain circumstances. More memory and processor speed enable us to deal with surfaces and compute interesting things. Finally, there are new applications, driven by other technologies and capabilities.

3D Surface Reconstruction

- Noise
- Occlusions
- Geometric Distortion (Calibration)
- Registration

“Registration and Reconstruction from the China Lake Range Data”, Whitaker and Juarez-Valdes, *University of Utah Technical Report*, 2001.



Data Courtesy U.S. Naval Air Warfare Center, China Lake, CA

Whitaker, ONR Sept. 19, 2003

An example of NavAir data shows the issues with measuring surfaces. This is airborne LADAR data taken from about 1km. It has noise (random and systematic) and it has occlusions, so that even “rectification” does not produce a full 3D model. A reconstruction system should combine several such scans from different views and create a full 3D model.


Research Strategy

- Establishing the relationships between measured data and surfaces
 - Statistical foundation–estimation theory
 - Estimating surfaces directly from ladar/sonar/radar
- New priors – methods for smoothing/processing surfaces
 - Generalizing image processing to surfaces
- Surface representations
 - Extending level set technology for surface modeling
 - New algorithms/computational mechanism

Whitaker, ONR Sept. 19, 2003

This is how we structured our research. Some theory relating surfaces to data. Some better ways of filtering or smoothing surfaces. And some new ways of representing surfaces and performing the associated computation.

Example Scenario



Whitaker, ONR Sept. 19, 2003

An example of how this type of capability could be used. Reconnaissance.


Major Results

- **Surface estimation**
 - Statistical formulation for surface reconstruction
 - Ladar and tomography
 - Freeform surfaces and parametric models
- **Surface processing and analysis**
 - Higher-order priors for surface reconstruction
 - General framework for mapping image processing algorithms to surface
- **Level-set surface modeling**
 - New algorithms for fast propagation/tracking of interfaces
 - New computational scheme for higher-order flows

Whitaker, ONR Sept. 19, 2003

Here are the results that came out of this project. All of these results have associated published papers and software.

Surface Estimation A Systems Approach




Bayesian reconstruction: maximizing the posterior – likelihood (data) + prior

Whitaker, ONR Sept. 19, 2003

This is the big picture of how to formulate the problem of “estimating” surfaces from noisy data. This is an extension of a philosophy which has been pursued with some success in the image processing literature. This is a Bayesian formulation where the surface is meant to optimize the statistical relationship between the data (likelihood) and what we know about shapes in the world (prior).

Surface Estimation from Ladar

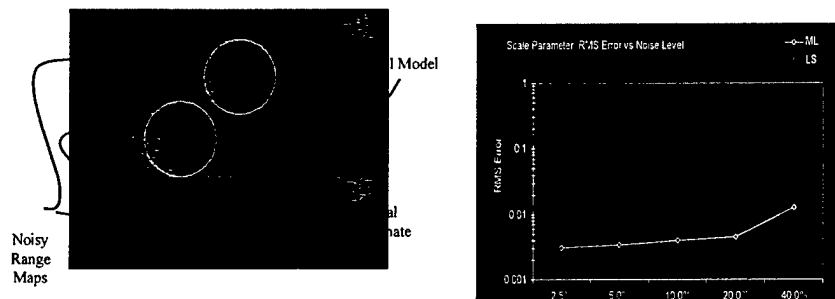
- Scanner Model
 - Geometry
 - Lines of sight
 - Pose
- Sensor Model
 - Noise
 - Gaussian w/outliers
- Shape Optimization
 - Parameters
 - Free-form



Whitaker, ONR Sept. 19, 2003

This is the “sensor model” it describes how, given a surface, a particular measurement would be produced. A typical ladar (laser range finder) takes a 2D array of measurements. Each measurement corresponds to a line of sight and gives a distance to the nearest surface, with some random noise. Studies of such instruments [Gregor and Whitaker, 2002] show that a Gaussian noise model with outliers is a good model. If we assume the measurements are independent, then we can quantify the “likelihood” of a particular surface, conditional on a set of range measurements.

Estimating 3D Surface Parameters Via Maximum Likelihood



- Converges more reliably and more accurately than ICP.
- Unbiased!

Whitaker, ONR Sept. 19, 2003


If we change the parameters of a surface (e.g. its orientation and size), we can try to find the surface that gives the best estimate of likelihood. This is a maximum likelihood estimator. Statistical theory says that this estimator is “unbiased”, which means on the average, it gives the correct answer. Using this estimator we can align models to range measurements. This slide shows some experiments in 2D to confirm that this is a reliable, accurate way to fit surfaces to range data. It does significantly better than the conventional approach (closest point error), which is biased.

Surface Estimation

Example: Ladar Data From Indoor Scene



Surface Rendering of Single Scan




Close Up View

Whitaker, ONR Sept. 19, 2003

This is some example range data. It is rendered as a surface. The triangles are constructed by connecting nearby points in the range image. This data is noisy. A closeup of the legs of the chair shows how little information is contained there.

Registration of Ladar Data


- Complex scenes w/noise
 - Improved accuracy and convergence



Whitaker, ONR Sept. 19, 2003

One of the big problems of constructing 3D models from range data is determining the relative position of the scans. Typically external approaches such as GPS or dead reckoning are not accurate enough to get good 3D reconstructions. Thus, people often rely on data-driven approaches for aligning different scans. We can use this maximum-likelihood approach for aligning models to data for registering one range map to another. The results proved to be very robust, even for complicated scenes such as that shown here (one scan rendered in white the other in magenta).


China Lake Data Registration



Whitaker, ONR Sept. 19, 2003

This an example of four different scans, taken from different point of view, registered into one coordinate system.


Surface Reconstruction



Whitaker, ONR Sept. 19, 2003

Once we have sets of registered scans, we can fit free-form surface models (using the same statistical approach) to multiples scans to get 3D models of complex scenes. In this case the models are represented volumetrically as level sets and the surface estimation includes a prior (to be describe more later) that prefers smooth surfaces with sharp creases. This reconstructed includes ten scans from different points of view.


Deltasphere LADAR Scan



Whitaker, ONR Sept. 19, 2003

This is data from the 3rd Tech Delta Sphere scanner.

Volumetric LADAR Reconstruction



Whitaker, ONR Sept. 19, 2003

A reconstruction from eight range images.


China Lake Reconstruction



Whitaker, ONR Sept. 19, 2003

Reconstruction from the China Lake data set.


China Lake Reconstruction



Whitaker, ONR Sept. 19, 2003

Higher-Order Surface Processing Feature-Preserving Flows

- Image Processing
 - E.g. anisotropic smoothing (Perona & Malik)



Whitaker, ONR Sept. 19, 2003

Part of the technology which makes this possible is a generalization of ideas from image processing. Using PDEs that penalized gradients of the image, people have been able construct nonlinear filters that smooth images and enhance strong edges. These terms can be used as “priors” and combined with a likelihood (data fitting) term to do reconstruction. One of our goals was to do this for surfaces.


Strategy

- Variational formulations of curvature
 - Generalization of grad mag from IP
 - Gradient of normal map
 - Allowance for outliers
- Decouple the normals from surface
 - Process normals
 - Refit surface to normal map
 - System of 2nd-order equations


Whitaker, ONR Sept. 19, 2003

The strategy for surfaces is to construct energy functionals that rely on penalizing the gradients of the normals, which live on the surface. If this is formulated correctly, the results will be “geometric”, which means they depend only on the shapes of the objects and not the parameterization or coordinate systems that one uses. The resulting PDEs are not easy to solve, because they include fourth-order derivatives. We have proposed a strategy that simultaneously smooths the normals of the surface while trying to keep the surface aligned with those normals.


Surface Filtering Strategy



Surface & Normals



Process Normals



Refit Normals


Whitaker, ONR Sept. 19, 2003

This is an example of how this computational strategy works. A shape induces a field of normals which are consistent with the surfaces nearby. We can smooth those normals (maintaining unit length). Then we can refit the surface. In this case a square would get “blurred” to a circle, which is a steady state of this process.

Anisotropic Diffusion



Original model

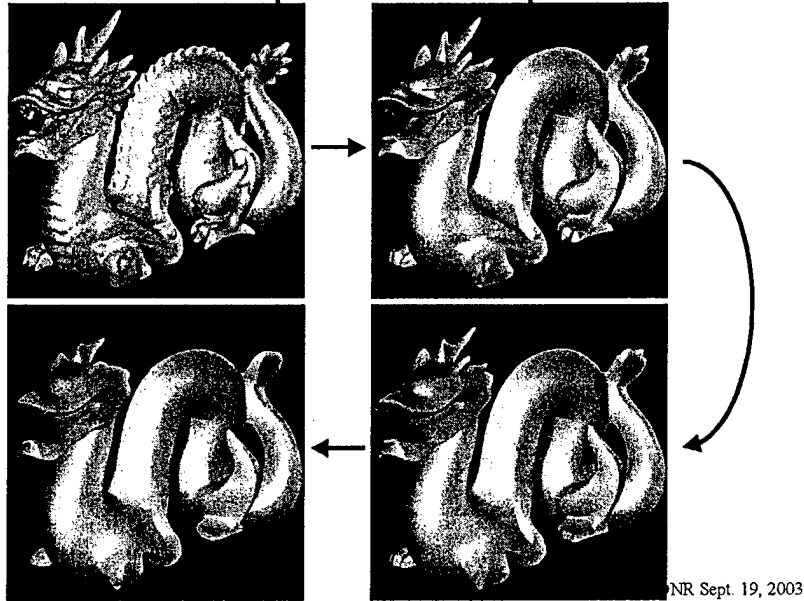


Anisotropic 4th order smoothing

Whitaker, ONR Sept. 19, 2003

If we use a penalty function on the gradient of the normals that allows for outliers, analogous to Perona&Malik for images, we obtain flows that smooth surfaces while enhancing strong creases.

Anisotropic Scale Space



Just like with images, this produces an anisotropic scale space, or a series of “cartoon-like” shapes which produces simpler objects that retain the important features.


Surface Analysis

- In the context of models for AR
- “It would be great if a soldier knew how many steps it would take to get from building A to vehicle B”
- Segmentation, hierarchical grouping, ...
- Low level -> high level

Whitaker, ONR Sept. 19, 2003

One of the goals of this work is to produce segmented surfaces that could be used for analysis. Thus, using generalizations of ideas from image processing, we can partition surfaces into meaningful pieces.


Curvature – Generalization of Edges



Whitaker, ONR Sept. 19, 2003

The “edges” of surface patches are places where the gradient of the normals (I.e. curvature) is high.


Watershed Segmentation



Whitaker, ONR Sept. 19, 2003

Using watershed segmentation we can produce a hierarchy of surface patches, which can be used to decompose a model into its constituent parts.

Watershed Segmentation



Whitaker, ONR Sept. 19, 2003



Watershed Segmentation



Whitaker, ONR Sept. 19, 2003

Cramer-Rao Bounds

A Statistical Estimate of Reconstruction Quality



Whitaker, ONR Sept. 19, 2003

Because we are using statistical formulations, we can rely on ideas from signal processing to compute error bounds. This is a generalization of Cramer-Rao bounds, which is typically used for putting error bounds on parameter estimates. In this case we have applied it to shape, to give us bounds on the position of the surface in the direction of the normal. This ties back into the sensor model. Thus, a particular range image, reduces the error of the surface in a way that depends on the noise in the measurements and the direction from which it is taken.

Processing Height Fields

- Formulation of higher-order smoothing for 2.5D surfaces
- Surface deformation gives PDE on image
- Appropriate for terrain processing
- Preprocessing + detection


Whitaker, ONR Sept. 19, 2003

Some applications done require full 3D surfaces, but rather height fields, in which one parameter is a function of the other two. This is useful, for instance, in dealing with terrain maps.

Terrain Analysis

Contour-Line Vertices – Ridges and Valleys


- Zero crossing of *third-order* differential quantities
- Noise and small-scale structure dominates
- Generally difficult to tune blurring, thresholds, etc.
- Too-many features + broken lines



Whitaker, ONR Sept. 19, 2003

For instance, one of the things people want to do with terrain data is detect features, such as valleys and ridges. This is useful, for instance, for compression. Valleys and ridges are indicated by high curvature in the contour lines. Detecting these requires third-order derivatives, which are noisy. Preprocessing can help this.

Anisotropic Diffusion for Preserving Ridges and Valleys



Whitaker, ONR Sept. 19, 2003

The same ideas used for level-set surfaces, which smooth data while preserving creases, apply to height fields, improving the detection of features.

Overcoming The Computational Burden Interactive, 3D, Deformable Models

- Nonlinear PDE – Iterative over volume
- Algorithms: front-tracking schemes
- Hardware:
 - Clusters and multiprocessors
 - Scalable
 - Commodity graphics cards
 - Price to performance
 - Rate of improvement – specialized processors

Whitaker, ONR Sept. 19, 2003


Many of the methods developed in this work require solutions to nonlinear PDEs, which can be computationally expensive. Part of the work is the development of methods for solving these problems more quickly. We have researched two approaches, one using multiprocessors on large super computers (e.g. SGI Origin series) and the other using commodity graphics cards (GPUs). Both approaches have been successful, but the GPU approach is more practical for a wider range of potential applications.

Solution: Graphics Processor

- Why?
 - Data level parallelism
 - High speed memory
 - Computation power going up quickly
 - Inexpensive (\$400 vs. \$50,000+)
- Issues
 - GPUs – SIMD processors, limited capabilities
 - Programming interfaces primitive
 - Limited GPU-CPU bandwidth/speed


Whitaker, ONR Sept. 19, 2003

GPUs are cheap and their power is increasing rapidly. They represent an example of a “streaming architecture”, which provides homogeneous, fine-grained parallelism (SIMD) with very high bandwidth to memory. Thus they are particularly well suited for solving PDEs on large grids.



Whitaker, ONR Sept. 19, 2003

This video shows a slice of a 3D model as it evolves on the GPU. The greyscales indicated curvature values and the little squares are “tiles” which allow us to focus the GPU processing only on those areas of the computational domain that are near the moving wavefront, which is associated with the level-set PDE.



This is an demo of an application for 3D segmentation using level-set based deformable models. The audio describes the specifics.

Publications

- R. Whitaker, "A level-set approach to 3D reconstruction from range data", *International Journal of Computer Vision*, 29(3), Oct. 1998, pp. 203–231.
- D. L. Elsner, R. Whitaker, and M. A. Abidi, "Volumetric modeling of objects and scenes using range images," *Digital Signal Processing: A Review Journal*, 9(2), April 1999, pp. 120–135.
- A. Mangan, R. Whitaker, "Partitioning 3D surface meshes using watershed segmentation", *IEEE Trans. on Visualization and Computer Graphics*, 5(4), Dec. 1999, pp. 308–321.
- R. Whitaker, "A level-set approach to image blending", *IEEE Trans. on Image Processing*, 9(11), Nov. 2000, pp. 1849–1861.
- D. Breen, R. Whitaker, "A level-set approach to 3D shape metamorphosis", *IEEE Trans. on Visualization and Computer Graphics*, 7(2), 2001, pp. 173–192.
- J. Gregor, R. Whitaker, "Indoor scene reconstruction for sets of noisy range images", *Graphical Models* 63(5), Sept. 2002, pp. 304–332.
- R. Whitaker, E. L. Valdes-Juarez, "On the reconstruction of height functions and terrain maps from dense range data", *IEEE Trans. on Image Processing*, 11(7), 2002, pp. 704–716.

Whitaker, ONR Sept. 19, 2003

Publications

- R. Whitaker, J. Gregor, "A maximum likelihood surface estimator for dense range data", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), October 2002, pp. 1372–1387.
- R. Whitaker, V. Elangovan, "A direct approach to estimating surfaces in tomographic data", *Journal of Medical Image Analysis*, 6(3), 2002, pp. 235–249.
- T. Tasdizen, R. Whitaker, P. Burchard, S. Osher, "Geometric surface processing via normal maps", To appear: *ACM Trans. on Graphics*.
- T. Tasdizen, R. Whitaker, "Higher-Order Geometric Surface Priors for Surface Estimation", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, To appear.
- A. Lefohn, J. Kniss, C. Hansen, R. Whitaker, "GPU-Based Interactive Level Sets for Volume Visualization and Analysis", *IEEE Trans. on Visualization and Computer Graphics*, To appear.

Whitaker, ONR Sept. 19, 2003

Publications (Conferences)

- R. Whitaker, D. Breen, "Level-set models for the deformation of solid objects," in *Proceedings of Implicit Surfaces '98*, Eurographics/Siggraph, June 1998, pp. 19–35.
- R. Whitaker, J. Gregor, and P. Chen, "Indoor scene reconstruction from sets of noisy range images", *Second International Conference on 3-D Digital Imaging and Modeling*, Oct. 1999, pp. 348–357.
- R. Whitaker, "Reducing aliasing artifacts in isosurfaces of binary volumes", *IEEE Volume Visualization And Graphics Symposium*, 2000, pp. 23–32.
- R. Whitaker, Xinwei Xue, "Variable-conductance, level-set curvature for image denoising", *IEEE International Conference on Image Processing*, October 2001, pp. 142–145.
- R. Whitaker, "Reconstructing terrain maps from dense range data", *IEEE International Conference on Image Processing*, October 2001, pp. 165–168.
- A. Gothandaraman, R. Whitaker, J. Gregor, "Total variation for the removal of blocking effects in DCT based encoding", *IEEE International Conference on Image Processing*, October 2001, pp. 455–458.

Whitaker, ONR Sept. 19, 2003

Publications (Conferences)

- T. Tasdizen, R. Whitaker, P. Burchard, S. Osher, "Geometric Surface Smoothing via Anisotropic Diffusion of Normals", *Proc. IEEE Visualization 2002*, pp. 125–132.
- K. Museth, D. Breen, R. Whitaker, A. Barr, "Level-Set Surface Editing Operators", in *SIGGRAPH 2002*, pp. 330–338.
- A. Lefohn, J. Kniss, C. Hansen, R. Whitaker, "Interactive GPU-Based Level-Set Models for Volume Visualization and Analysis", *Proc. IEEE Visualization 2003*, To appear.
- A. Lefohn, J. Cates, R. Whitaker, "GPU-Based Level-Set Models for Brain Tumor Segmentation", *Proc. Medical Image Computing and Computer Aided Intervention 2003*, To appear.
- T. Tasdizen, R. Whitaker, "Cramer-Rao Bounds on Surface Estimation", *Proc. 3Dim 2003*, To appear.
- T. Tasdizen, R. Whitaker, "Higher-Order, Feature-Preserving Priors for Surface Estimation", *Proc. 3Dim 2003*, To appear.

Whitaker, ONR Sept. 19, 2003

Tech Transfer – Software

- **VISPack**
 - C++, object-oriented toolkit
 - Image/volume processing, level-set solvers
 - Several hundred downloads, mostly academics
- **Insight Toolkit**
 - NIH/NLM Sponsored
 - First release this fall – thousands of downloads/month
 - Ported most of the level-set and reconstruction algorithms developed under ONR project
 - E.g. recent contract for fourth-order surface flows

Whitaker, ONR Sept. 19, 2003

Workshops That Presented This Work

- "Beyond Blobs" SIGGRAPH, 2003
- "Image Processing for Volume Graphics" SIGGRAPH 2002
- "Beyond Blobs" SIGGRAPH, 2002
- "PDEs for Graphics and Image Processing" SIGGRAPH, 2002
- "Image Processing for Volume Graphics" IEEE Visualization, 2002
- "Image Processing for Volume Graphics" SIGGRAPH 2001
- "Image Processing for Volume Graphics" IEEE Visualization, 2001

Whitaker, ONR Sept. 19, 2003

Students and Post-Docs

- Alan Mangan
- Xinwei Xue
- Ernesto Juarez
- Suyash Awate
- Vidya Elangovan
- Tolga Tasdizen

Whitaker, ONR Sept. 19, 2003