Continuous Terrain Modeling from Image Sequences with Applications to Change Detection

Yvan G. Leclerc[*]
Artificial Intelligence Center, SRI International
333 Ravenswood Ave., Menlo Park, CA 94025
E-MAIL: leclerc@ai.sri.com
HOMEPAGE: http://www.ai.sri.com/~leclerc/
PROJECT HOMEPAGE: http://www.ai.sri.com/~meshes/

Abstract

The objective of this project is to develop methods to incrementally model and detect changes in the shape or surface material properties of terrain. The model will be derived from range data (such as interferometric synthetic aperture radar (IFSAR) elevation data) and electro-optical (EO) and infrared (IR) imagery. The imagery can be either still images or video data, such as that obtained from the Predator UAVs. The modeling and change detection algorithms will be based on an extension of our object–centered "deformable mesh" approach that incorporates surface material properties and appropriate error estimates.

1 Introduction

The deployment of various monitoring platforms, such as the Predator UAV, will generate large quantities of SAR/IFSAR and EO/IR data of great value to Battlefield Awareness if it can be interpreted quickly and affordably.

We propose to develop and demonstrate a system that will automatically generate and refine a 3-D model of the terrain's shape and surface properties from IFSAR, EO, and IR data, and detect changes in both elevation (due perhaps to bomb damage, movement of large machinery, deforestation, and so on) and surface properties (due perhaps to change in ground cover, pouring asphalt over a dirt road, building an air strip, and so on). Such changes can then be noted on the model for review and action. We thus expect to be able to dramatically reduce the amount of analyst time necessary to take advantage of this type of data.

2 Overview of Our Approach

We propose to use our object-centered "deformable mesh" representation to combine radar and EO/IR imagery taken at different times of day and from different points of view into a unified 3-D model of the shape and surface properties of the terrain [Fua and Leclerc, 1996, Fua and Leclerc, 1995, Fua and Leclerc, 1994a, Fua and Leclerc, 1994b, Fua and Leclerc, 1993].

In this approach, the terrain is represented by a 3-D surface model composed of interconnected triangles called a "mesh." Each triangle, or facet, of the mesh represents an estimate of the position, shape, orientation, and surface material properties (e.g., color, radar reflectance) of the terrain's surface over a small triangular area. The mesh is used not only as the representation of the terrain, but is also an integral

---

[*] This work was sponsored in part by the Defense Advanced Research Projects Agency under contract F33615-97-C-1023 monitored by Wright Laboratory. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the United States Government, or SRI International. Many thanks to Pascal Fua for the many years of collaboration we have had in developing the approach described here.

part of the computational framework. Figure 1 illustrates the mesh representation and shows a mesh constructed from a stereo pair of EO images.

The method involves the following components.

Model Creation First, maps, IFSAR elevation data, EO/IR imagery, and other sources of information (e.g., terrain type, building models) will be combined to create a complete 3-D model of the shape and surface properties of the geographical area covered by the data. The surface properties of the terrain will be estimated from the imagery based on the known position of the sun, the shape of the terrain (taking shadows and occlusions into account), camera and radar parameters, cloud cover, and other relevant information. Known "deficiencies" of the sensors (such as occluded areas in EO/IR imagery or "front-porch" artifacts in IFSAR data) will be used to rigorously derive error tolerances and covariances for every element of the model. (In the later years of this effort, we expect that SAR data will be directly integrated into the model.)

Change Detection Second, new imagery and new IFSAR elevation data will be compared against the terrain model to detect changes in the terrain. It is the integrated 3-D nature of our representation and processing methodology that will allow us to detect changes in both the shape and surface material properties of the terrain, as follows.

Changes in the terrain's shape will be detected by comparing 3-D shape and material properties derived from incoming data against the model. This can be done directly for incoming IFSAR range data. For incoming EO/IR imagery, our mesh-based terrain modeling algorithm will be used to register and derive a new 3-D model from the imagery. This derived model will then be compared against the current model, using the error tolerances mentioned above to detect areas of significant change.

Model Refinement Third, new imagery and elevation data will be used to continuously refine the terrain model wherever the new imagery is consistent with the model (i.e., when the elevation data and surface properties derived from the new imagery are within the automatically derived error tolerance of the model). This will allow the model to become increasingly accurate and reliable over time. As the model becomes more accurate, it will support more sensitive change detection.

Model Extension Finally, incremental extensions of the model to new areas will be made wherever IFSAR range data or overlapping images cover a portion of the terrain that has not yet been modeled.

In the following sections we describe our approach and the proposed processing steps in more detail.

3 Mesh-Based Optimization

In mesh-based optimization, information from elevation data and imagery is integrated using a unified optimization framework in which a global objective function is minimized. Each source of information is modeled using a distinct objective function that relates the information to the shape and surface material properties of the surface mesh. A weighted sum of the objective functions is minimized to arrive at a model that incorporates all the information.

We propose to extend our current mesh-based approach in two ways so that it can be used for long-term model building and change detection based on IFSAR range data and sequential imagery. First, we propose to include rigorously derived error tolerances and covariance matrices that specify the range of positions/surface properties of each facet. Second, we propose to modify the optimization procedure so that new data will be processed in sequence as it arrives, rather than waiting for all the data to arrive before processing it.

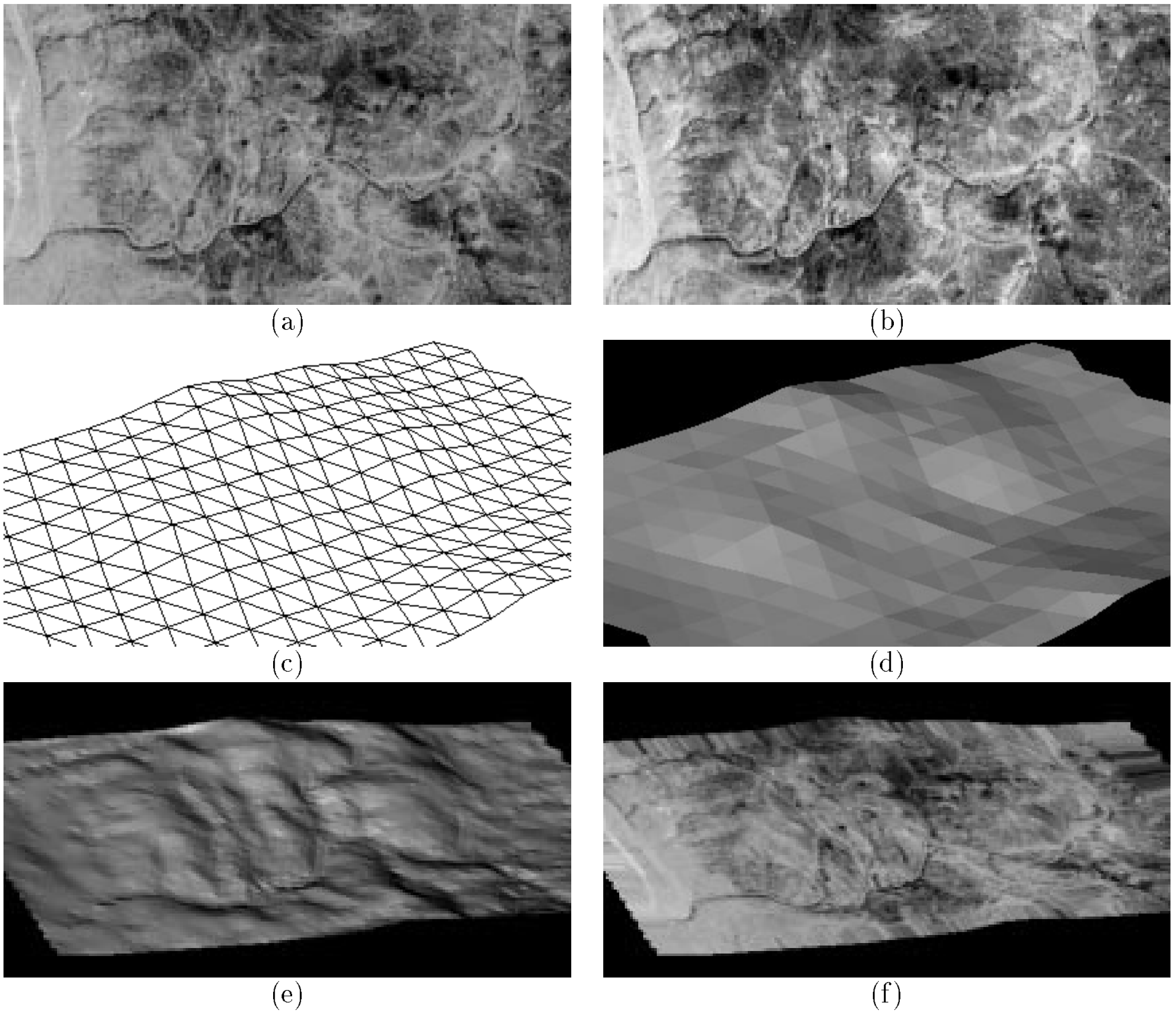

Figure 1: Terrain modeling at the Ft. Irwin, CA, National Training Center (NTC) using deformable meshes. (a,b) A stereo pair of a hilly site. (c) A coarse hexagonally triangulated mesh, shown as a wireframe. (d) A shaded view of the same mesh. (e) The mesh after subdivision and optimization, shown as a shaded surface. (f) The optimized mesh shown with one of the images overlaid on the surface.

3.1 Measures of Uncertainty

We propose to include a rigorously derived error tolerance for every element of the mesh that specifies the range of positions and surface material properties that is consistent with the quantity and quality of data processed to date. For example, areas that are covered by many views would generally have a smaller error tolerance than areas that have been viewed only a few times. Another example is that the elevation in an IFSAR shadow area is not well defined: the minimum elevation is unconstrained, but the maximum elevation can be computed from the look-angle. These constraints form the error tolerances for the shadowed area. Error tolerances can also be determined from other information sources that can be used to augment the model-building process, such as maps or annotations associated with the IFSAR range data. The rigorous derivation and use of error tolerances will be a significant component of our research effort.

In addition to the error tolerance, each element of the mesh has an associated covariance and information matrix. The covariance matrix represents the degree of uncertainty in the element's state. It is closely related to the error tolerance, and is directly related to the shape of the potential surface (defined by our objective function) in the neighborhood of the current estimate. The information matrix is the sum of the inverse covariances of the data used to update the element. As new data arrives, its information matrix is added to the information matrix of the associated model element.

3.2 Model Initialization

An initial terrain model must be created before refinement and change detection can take place. If a reasonable estimate of the shape of the terrain is provided (from IFSAR range data or maps, for example), then the initial model creation can be done automatically. Otherwise, manual intervention will be required. For more information on how manually derived information is used to initialize a mesh, see [Fua and Leclerc, 1995].

3.3 Data Processing

Once an initial model is created, incoming IFSAR range data and imagery will be used in three ways: to detect and correct for errors in the model, to detect changes in the area, and to update the terrain model to make it more accurate.

Incoming imagery will be processed as follows. Incoming imagery (both elevation and EO) is used to create an updated terrain model. This updated model is compared against the current model and corresponding error tolerances. Areas that fall outside of the error tolerances are candidates for model correction or denote changes to the scene. Areas that are inside the error tolerance, on the other hand, are used to refine the model with the optimization process.

For IFSAR range data, the comparison is relatively straightforward. Every point of the updated elevation data is directly compared against the current surface mesh. If it is within the error tolerance, then that part of the data can be used to refine the surface mesh by using a standard Kalman filtering approach. In that approach, the coordinates of the model element are replaced by a weighted average of the model point and the new elevation data, where the weights are the information matrices described earlier.

Isolated points or small areas that fall outside of the error tolerances are likely to be sensor errors that will be ignored. Large areas that fall outside of the error tolerances indicate that either the terrain has changed, or that the mesh is in error. One way to distinguish between these two cases is to re-optimize the mesh using previous imagery. If the mesh changes (that is, the previous imagery is incompatible with the new model, and hence is incompatible with the new imagery), then this indicates that the terrain has changed.

Comparison of EO imagery is more complex. The scenario we envision for incoming imagery is that the images will be processed as sequences in which adjacent frames are taken relatively close together (such as would be the case for UAV video data of a continuous fly-over). A large gap (in time) between adjacent frames will be treated as the beginning of a new sequence, and the remaining frames will be processed as a separate sequence.

Sequences of EO images, as defined above, can be used to refine or detect changes in either the shape or the surface material properties of the terrain. In all cases, the basic idea is to use the sequence to estimate an updated terrain model and then compare this updated terrain model against the current model.

The updated terrain model will be computed by starting with the current model as the initial estimate of our optimization procedure. Each new frame of the incoming imagery will then be used to refine this updated model, using a sequential version of our optimization procedure. Eventually, the covariances of the updated model will become small enough to allow a meaningful comparison against the current model. Since the updated model will contain both the shape and surface reflectance properties, it will be possible to detect changes in both of these aspects of the terrain.

When the sequence is complete, the updated and current terrain models will then be merged (and the covariances appropriately adjusted) wherever the differences are within the error tolerances. Note that, over time, areas viewed multiple times will tend to have lower covariances. Consequently, the change detection will become more sensitive over time.

4 Advantages of the Mesh Approach

Deformable meshes have a number of distinct advantages over traditional image-based stereo and change detection techniques.

- Occlusions in arbitrary views are naturally accommodated because meshes are a full 3-D representation of a terrain. Traditional stereo techniques, on the other hand, require that occlusions be detected explicitly in the images, which is an open research problem.
- Information from many modalities can be naturally integrated within the unified optimization framework. This approach produces a model in which all the information is used together. This is significantly more accurate and robust than traditional processing where, for example, independent depth maps are recovered from stereo pairs and the maps are then "averaged" together in some fashion.
- Constraints from various external sources, such as maps, can be incorporated by creating an appropriate objective function or by constraining the optimization process in the relevant manner.
- The expected accuracy and known artifacts of various sensors can be incorporated into the modeling process by appropriately weighting components of the objective functions. For example, "shadowed" areas in IFSAR data would be weighted very lightly, while data from flatter areas would be weighted more heavily. In addition, error tolerances and covariance matrices specifying the range of positions for the surface elements can be derived from the expected accuracy of the sensors.
- Change detection using new images can be accomplished even when viewpoints and time of day have changed because the terrain model is fully 3-D and includes surface properties. Change detection based on a simple comparison of images (such as current mosaicking techniques) cannot be used for this purpose.

5 Evaluation Plan

We will provide metrics to evaluate the accuracy, robustness, and completeness of the terrain models we produce, as well as the robustness and accuracy of the change detection.

- **Accuracy of the model.** The accuracy of the terrain model will be measured against a number of standards: points on the terrain with known positions (as obtained via Global Positioning System (GPS) sensors on the ground), selected points in images for which the best manual photogrammetry has been applied, carefully monitored automatic stereo analysis systems, and IFSAR elevation data in areas for which IFSAR data had not been supplied to the system.
- **Robustness of the model.** Robustness will be measured in terms of the area of the recovered terrain in which the system made clear mistakes (again as compared to human-recovered terrain models).
- **Completeness of the model.** Completeness will be measured in terms of the area of the recovered terrain for which the system had at least two views but that was not modeled.
- **Accuracy and robustness of the change detection.** The accuracy and robustness of the change detection will be measured by the number of missed changes and the number of false positives generated by the algorithm.

6 Summary

In summary, we propose to develop methods to incrementally model and detect changes in the shape or surface material properties of terrain. This will be done by extending our current mesh-based optimization approach in two ways, so that it can be used for long-term model building and change detection based on IFSAR range data and sequential imagery. First, we propose to include rigorously derived error tolerances and covariance matrices that specify the range of positions/surface properties of each facet. Second, we propose to modify the optimization procedure so that new data will be processed in sequence as it arrives, rather than waiting for all the data to arrive before processing it.

We have proposed a number of methods for robustly detecting changes in 3–D meshes using our approach. These proposed methods are still in the preliminary stages of development, and we will certainly be considering other recent work in change detection to see if some of the techniques can be applied to 3–D meshes [Huertas and Nevatia, 1996, Bejanin *et al.*, 1994, Chellappa *et al.*, 1994].

References

[Bejanin *et al.*, 1994] M. Bejanin, A. Huertas, G. Medioni, and R. Nevatia. Model validation for change detection. In *Proceedings of the DARPA Image Understanding Workshop*, pages 287–294, Monterey, CA, November 1994.

[Chellappa *et al.*, 1994] R. Chellappa, Q. Zheng, L. S. Davis, C. L. Lin, X. Zhang, C. Rodriguez, A. Rosenfeld, and T. Moore. Site-model-based monitoring of aerial images. In *Proceedings of the DARPA Image Understanding Workshop*, pages 295–318, Monterey, CA, November 1994.

[Fua and Leclerc, 1993] P. Fua and Y. G. Leclerc. Combining Stereo, Shading and Geometric Constraints for Surface Reconstruction from Multiple Views. In *SPIE Workshop on Geometric Methods in Computer Vision*, pages 113–123, San Diego, CA, July 1993.

[Fua and Leclerc, 1994a] P. Fua and Y. G. Leclerc. Registration Without Correspondences. In *Conference on Computer Vision and Pattern Recognition*, pages 121–128, Seattle, WA, June 1994.

[Fua and Leclerc, 1994b] P. Fua and Y. G. Leclerc. Using 3–Dimensional Meshes To Combine Image-Based and Geometry-Based Constraints. In *European Conference on Computer Vision*, pages 281–291, Stockholm, Sweden, May 1994.

[Fua and Leclerc, 1995] P. Fua and Y. G. Leclerc. Object-Centered Surface Reconstruction: Combining Multi-Image Stereo and Shading. *International Journal of Computer Vision*, 16:35–56, September 1995.

[Fua and Leclerc, 1996] P. Fua and Y. G. Leclerc. Taking Advantage of Image-Based and Geometry-Based Constraints to Recover 3–D Surfaces. *Computer Vision and Image Understanding*, 64(1):111–127, July 1996. Also available as Tech Note 536, Artificial Intelligence Center, SRI International.

[Huertas and Nevatia, 1996] A. Huertas and R. Nevatia. Detecting changes in aerial views of man-made structures. In *Proceedings of the DARPA Image Understanding Workshop*, pages 381–388, Palm Springs, CA, February 1996.