Daedalus Battlefield Visualization System[1]

Edward Riseman, Howard Schultz, W. Jonathon Lim,
Bruce Draper, Allen Hanson
Computer Vision Laboratory
Computer Science Department
University of Massachusetts
Amherst, MA 0l003
URL: http://vis-www.cs.umass.edu/

Abstract

This paper describes the conceptual structure of the Daedalus battlefield awareness system. The motivation for Daedalus arises from the problem of providing timely situation awareness for air and ground combat operations in order to improve force effectiveness and reduce fratricide. Daedalus is being designed to fuse the flood of information from sensor platforms currently being developed into a single, consistent 3D model of the dynamic battlefield. Real-time rendering of terrain views overlaid with reconnaisance information and military symbology provides battlefield commanders with a powerful new tool for the planning of operations, an extremely effective interface for briefing the plan to those who will implement it, and a system for monitoring battlefield operations which incorporates up-to-date information.

I. Introduction

The problem of providing timely situation awareness for air and ground combat operations is a recurring theme that impacts both force effectiveness and the need to reduce fratricide. Typically, the resolution of terrain maps that are available for a new site of interest are at best 30 meter DTED (digital terrain elevation data) which may have been produced years before actual conflict. Similarly, the integration of new, possibly continuous, live information has been problematic due to communication bandwidth constraints and the dangers to human scouts. Consequently, current methodology for both the planning and the conduct of combat missions has rarely employed the kind of up-to-date information that is becoming possible via new technologies.

In contrast, the introduction of new remote sensor platforms for military reconnaissance will provide unprecedented coverage of the evolving battlefield. For example, the UAV Tier 2+ and Tier 3- programs (Delaney 1995) are expected to deploy reconnaissance aircraft which have the ability to stay aloft for over 24 hours and continuously monitor 40,000 sq. nm per day using visible light and infrared, and eventually synthetic aperture radar sensors. These aircraft are expected to produce datastreams of 50 Mbits per second to satellites and 137 Mbits per second to dedicated ground stations. Analogously, scenarios from the ARPA UGV Demo II program will involve up to four UGV scouts carrying out ATR missions from multiple, possibly overlapping, viewpoints.

It is obvious that novel means are required to provide commanders the ability to instantly comprehend dynamic battlefield situations and to react accordingly. The problem of battlefield awareness for the future commander is that he typically will not be co-located with the multiple, distributed sensors of future systems, and usually will not be intimately familiar with the terrain. The result can be a serious loss in situation awareness. The problem of information overload will be tremendously magnified by the future scenarios being developed, where much larger numbers of ground sensors are deployed by UGVs or other means, each capturing a local view that somehow must be integrated into a coherent and consistent up-to-date world view.

The Daedalus system, being developed for the Army Research Labs (ARL), is a first step toward applications that can take advantage of these resources, while avoiding the inherent problems of information overload and the confusion and chaos that could otherwise result. The system is designed to integrate sensor data from multiple remote sensors onto reconstructed battlefield terrain, yielding a single, consistent 3D model of the battlefield. The terrain map will be produced from

---

[1]This research has been supported in part by the Defense Advanced Research Projects Agency under contract DAAL02-91-K-0047 (through ARL); by the National Science Foundation under grant IRI-9208920; and by the Office of Naval Research under grants N00014-95-1-1035 and N00014-89-J-3229.

aerial images by the UMass terrain reconstruction system (Schultz 1994; Schultz 1995); this system is designed to produce accurate elevation maps from images that are possibly taken from highly oblique viewing angles, from widely separated viewpoints, and from sensors with large base-to-height ratios[2]. The goal is to achieve digital elevation maps (DEM) resolutions of about 1 meter on the ground. The UMass terrain reconstruction system has been qualitatively evaluated on data from the ARPA UGV sites at Lockheed-Martin and at Fort Hood; it has been shown to produce detailed terrain maps under widely varying viewing scenarios.

Once a DEM has been produced, it can be texture mapped with one of the images by standard techniques to produce a terrain map for a visualization database. ARL has been developing a visualization system in conjunction with Georgia Institute of Technology. The system, called VGIS (Virtual Geographical Information System) (Ribarsky, Op den Bosch et al. 1994), allows simulated fly-throughs in real-time or close to real-time. As part of the development of Daedalus, these terrain maps have already been interfaced to the ARL VGIS system.

As combat vehicles come into the range of UGV scouts or other sensors, their location, direction, velocity, and IFF status will be overlaid onto the terrain map producing a rendition of the live battlefield. ARL CAD models of military vehicles have been used in the initial simulations. These models were placed on the terrain at appropriate locations; both the CAD models and the terrain were then rendered in real-time in VGIS and a fly-through scenario was developed.

In the command post the system will allow the commander to freely roam the battlefield (viewing the situation from any perspective), without being restricted by the original sensor positions. Because the information about vehicle movements is transformed onto the rendered terrain database in an integrated and consistent manner, confusion regarding multiple viewpoints will be avoided - despite asynchronous information updates from multiple UGV scout and other ground sensors. Furthermore, the commander can, at any time, ask for more detailed information about battlefield conditions; in particular, image chips associated with particular vehicles that were derived from the deployed sensors can be viewed. An overview of the Daedalus system concept is depicted in Figure 1.

2. Overview of the UMass Terrain Reconstruction System

The UMass terrain reconstruction system (Schultz 1994) employs a hierarchical image matching scheme that is specifically designed to recover the three-dimensional structure of textured surfaces (e.g. ground elevation maps) from two overlapping images that *may* have been taken from oblique and widely separated viewpoints and with large base-to-height ratios.

When digital terrain maps (DTMs) are computed from widely separated images, perspective effects may result in a large number of false matches and poor reconstruction accuracy. When the base-to-height ratio (*b/h*) exceeds 0.5, the performance of image matching algorithms begins to deteriorate. When *b/h* becomes greater than 1, elevation errors caused by perspective can become quite large. Different views of the same ground patch at 45 degrees or more from nadir will produce highly foreshortened pixels on opposite sides of the patch, which reduces the ability of standard correlation algorithms to match correctly. For operational scenarios for UGV, RADIUS, and UAV, aerial images often are taken from distant reconnaissance sources, and an ability to deal with large base-to-height ratios would be advantageous.

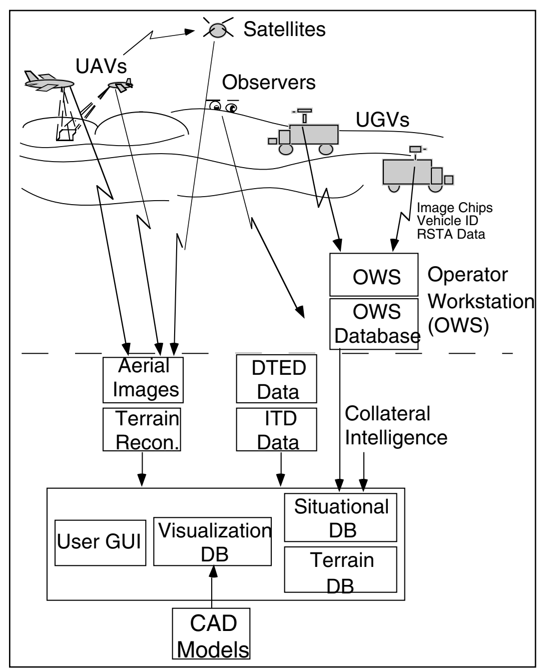

Figure 1. An Overview of DAEDALUS. Information from many sources is fused into a dynamic graphical representation of the battlefield situation.

[2] The base to height ratio is defined as the distance between camera centers (camera baseline) divided by the distance from the midpoint of the camera baseline to a nominal point on the surface.

When digital terrain maps (DTMs) are computed from widely separated images, perspective effects may result in a large number of false matches and poor reconstruction accuracy. When the base-to-height ratio (*b/h*) exceeds 0.5, the performance of image matching algorithms begins to deteriorate. When *b/h* becomes greater than 1, elevation errors caused by perspective can become quite large. Different views of the same ground patch at 45 degrees or more from nadir will produce highly foreshortened pixels on opposite sides of the patch, which reduces the ability of standard correlation algorithms to match correctly. For operational scenarios for UGV, RADIUS, and UAV, aerial images often are taken from distant reconnaissance sources, and an ability to deal with large base-to-height ratios would be advantageous.

The three key elements of the algorithm described below are modifications to standard approaches to stereo algorithms and are designed specifically to recover a DEM from images taken from widely varying viewpoints. These techniques include a new correlation match score, and modifications to standard subpixel and hierarchical matching techniques.

When image matching techniques are used to compute DEMs, a dilemma occurs when the size of the correlation mask is selected. To increase robustness to random noise the mask should be as large as possible, and to minimize the effects of perspective distortion the mask should be as small as possible. To develop algorithms that balance these competing factors we take advantage of the fact that pixels near the center of the correlation mask are less affected by perspective distortion. In many conventional methods, each element in the mask contributes equally to the overall match score. However, by assigning position-dependent weights, it is possible to place more emphasis on the central elements. A weight is assigned to each element in the correlation mask that depends on its distance from the mask center. Gaussian weights are used in which the standard deviation of the weights (σ) is either fixed or context dependent.

In addition, instead of estimating the optimal match score at integer disparity steps, a series of match scores are computed at subpixel disparity steps. By estimating the shape of the disparity function over a narrower disparity range at the last step in the correlation process, the peak of the disparity function is less affected by perspective distortion and hence more accurately defined.

Finally, an enhancement to standard multiresolution matching is employed to allow the system to converge on a disparity map which has a large range of expected disparities. Figure 2 shows the hierarchical control loop of the algorithm. Initially, multi-resolution image pyramids are generated for the two images using Gaussian reduction. Then, starting at the coarsest resolution level, a disparity map is estimated between the reference left image and the "warped" right image. This disparity map is used to remove viewpoint-induced distortions in the right image (relative to the reference image) by unwarping it using a resampling technique. An incremental disparity map at a finer resolution is then computed by matching the resampled warped image at the next finer resolution with the reference image at that level. This process of successively removing perspective distortion and computing an incremental disparity map is repeated until the reference image is matched to the warped image at the finest resolution level. The final disparity map is found by simply summing the incremental disparity maps computed at each resolution level.

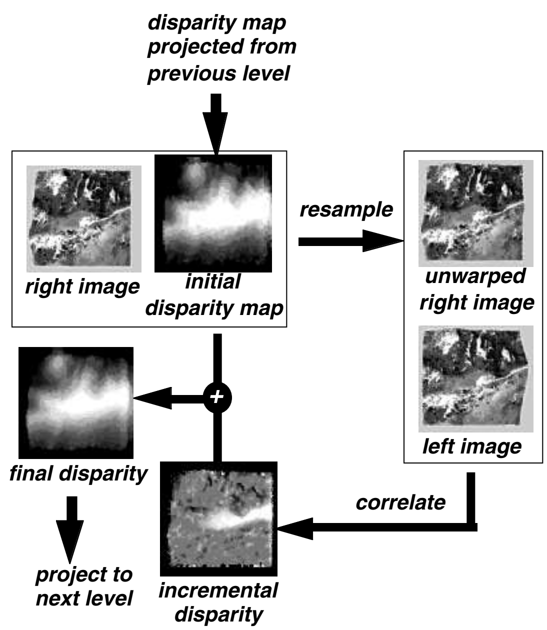

Figure 2. Multiresolution contruction of disparity maps. An initial disparity map is constructed from a reduced resolution pair of images. This map is projected to the next finer resolution level and used as a seed for the computation of a disparity map at this level. The process is repreated until the base image resolution is reached.

These several variations have been incorporated into a stereo analysis package which was carefully engineered on sound photogrammetric principles. We have compared the UMass terrain reconstruction system to standard techniques by processing sequences of simulated images with values of *b/h* that varied between 0.25 and 2.25, and a single pair of high altitude images with a base-to-height ratio of 0.63. Analysis of the simulated data showed that the reconstruction accuracy of the system remains independent of base-to-height ratio.

The system has also been applied to a series of real aerial image data over the past two years, including Ft. Hood and the Martin-Marietta Demo B site. Figure 3 shows two oblique views of Fort Hood taken at approximately 36 and 53 degrees tilt (off-nadir) and very different viewpoints.. The base-to-height ratio is 1.525. Also shown is a perspective view generated from the reconstructed DEM and orthoimage; the ground sample distance on the reconstruction is 1 meter. One of the two nadir views of the Martin-Marietta Demo B site is shown in Figure 4a, along with a reconstruction of a portion of the terrain without the large building. Additional data from these reconstructions will be shown in some of the next sections.

Two extensions to the terrain reconstruction system are being developed. The first is an automated match point (image correspondence) selection system for those cases when the pose information associated with the sensors is only approximately correct. One approach being examined is the use of "interest operators" based on deformable templates (Ravela, Draper et al. 1995) that develop a peaked autocorrelation response, resulting in the selection of distinguished image points that have high matchability. The second extension will allow the system to use multiple images (instead of just two) during the reconstruction process. The greatest improvement in terrain resolution and accuracy will come with the use of just a third image, since any ambiguity in a decision concerning the best match between two features almost always will be resolved by epipolar constraints from the third image.

3. Enhancement of the UGV Mission

Currently, standard terrain maps, for example 30 meter DTED, are used for mission planning. However, the terrain reconstructed using the algorithm described in the previous section is expected to be at or near 1-meter resolution. In addition, the DEM is texture-mapped so that when rendered from oblique views, it often will provide useful predictions of scene appearance. There are numerous ways in which this combined data can be used to advantage in the UGV program; some of these are briefly described below.

3.1. Mission Planning

In the UGV Demo II program, planning for scout vehicles takes place as a function of expected enemy locations and potential overlook positions. Up-to-date terrain maps with 1-meter resolution contain a much more detailed view of terrain characteristics. Using this information, it is possible to develop more accurate navigation plans and to estimate mission parameters, such as vehicle speed, using vehicle simulators operating on the specific terrain.

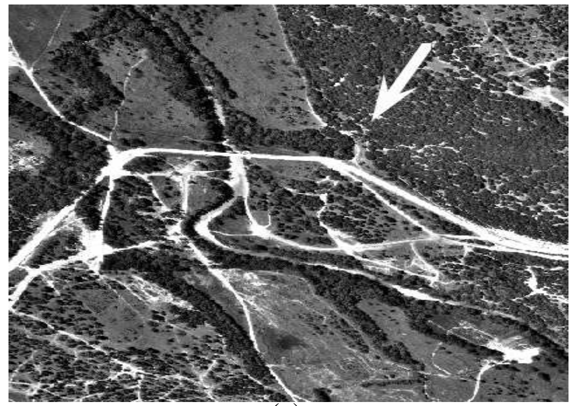

(a)

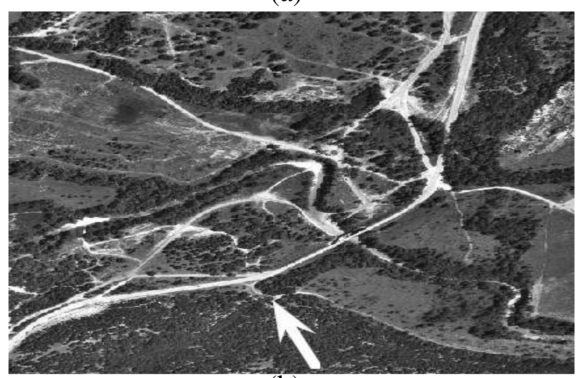

(b)

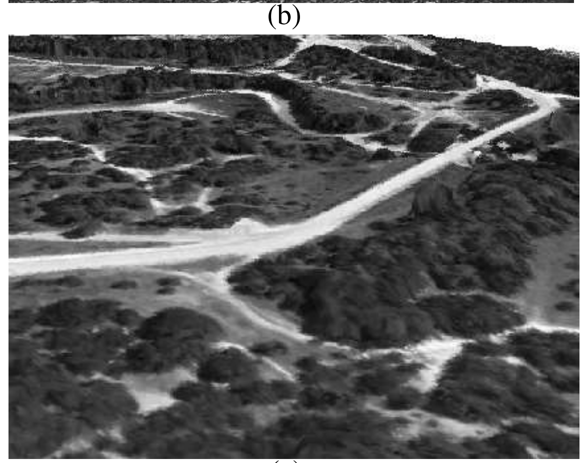

(c)

Figure 3. Fort Hood Terrain Reconstruction. (a,b) Two oblique views of Fort Hood taken at approximately 36 and 53 degrees tilt (off-nadir) and very different viewpoints. The base-to-height ratio is 1.525. (c) A perspective view generated from the reconstructed DEM and orthoimage. The ground sample distance on the reconstruction is 1 meter.

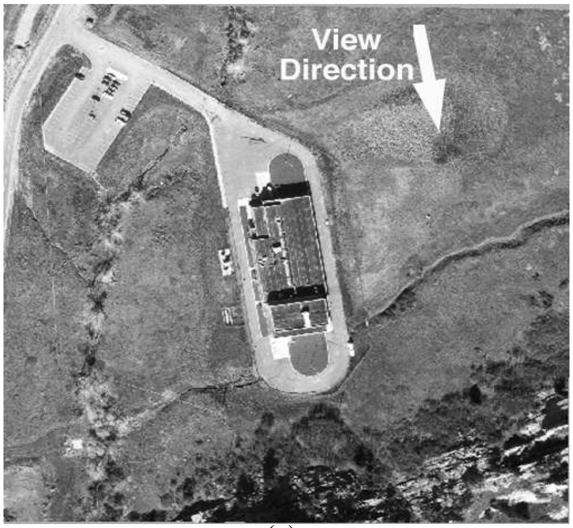

(a)

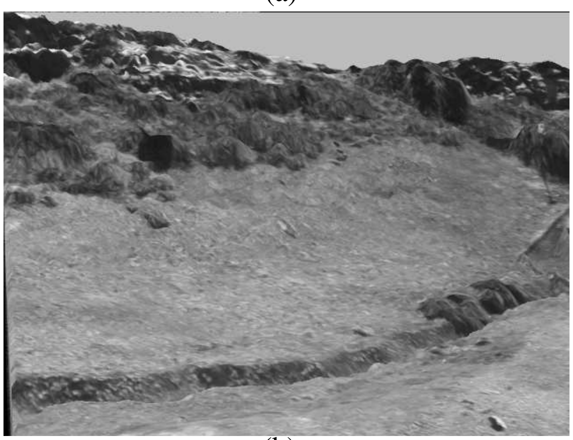

(b)

Figure 4. Lockheed-Martin/Denver Terrain Reconstruction. (a) One of the two nadir images used to generate the DEM and ortho-image; the white arrow shows the view direction for the portion of the reconstruction shown in (b).

If potential areas of enemy troops are hypothesized, "stealth planning" is greatly aided by such high resolution maps. Visibility analysis can be applied to the terrain map so that UGV paths will remain hidden from selected target regions, and if that is not possible, then the distance (or alternatively time) for which they would be visible can be minimized. Consider Figure 5 where a set of bounding overwatch positions is determined based on terrain characteristics, and a stealth path has been constructed from the current vehicle location to a point on the visibility boundary (Ravela, Weiss et al. 1994).

A variation on this theme is the use of the terrain and the visibility analysis during a RSTA mission in order to accurately position a vehicle to perform ATR. Here, the goal is to minimize the exposure of the vehicle to the enemy position. A maneuver called "hill cresting" is possible in which the terrain above or beyond the goal target location is used to servo the UGV, so that ideally only the sensor is visible and most of the vehicle is still hidden by the hill that has just been traversed; see Figure 6 for an illustrative example.

3.2 Obstacle Prediction/Detection

Currently UGVs must determine all obstacles via stereo or active laser ranging systems, or a combination of the two. With high resolution terrain maps, many vehicle obstacles arising from terrain characteristics can be anticipated. This ability effectively changes obstacle detection from a bottom-up discovery process to a top-down predict-and-verify process. In many cases rendering of ground level views may allow images of the actual obstacle to be presented a priori and speed up or improve obstacle recognition, avoiding the need for an active sensor.

As an example consider the shallow (1-1.5 meter deep and 1.5-2 meters wide) erosion ditch and associated elevation data shown in Fig. 7(a). The ditch is also shown in the Lockheed-Marietta terrain reconstruction (Fig 4(b))**.** The structure of the ditch is clearly visible. In addition, Fig 7(b) shows a close up view of the rock located in the bottom right corner of Fig 4(b). Fig 7(c) shows the elevation and gray scale profile along the dashed line shown in Fig. 7(b). In this example, the terrain map clearly provides enough information to detect obstacles as small as approximately 0.25 meters.

3.3. Pre-mission Automatic Target Detection (ATD)

The Daedalus terrain database for visualization represents a synthesis of optical data and a DEM. Thus, mechanisms can be developed as part of the reconstruction process to classify objects since the 3D representation augments at least two views of an object.

While typical ATD is applied to FLIR, SAR, and possibly visible data, these are typically unfused 2D representations of information. Here we have fused multiple representations (2D and 3D) of the same data. to which the 'standard' ATD techniques can be applied and the results fused in a natural way. Since the system would be used as a focus-of-attention mechanism, determining potential target locations for investigation by scout vehicles, it would not have to produce final classification decisions.

As an example consider the reconstruction of vehicles in the parking lot. The tops of vehicles in Fig. 8 are on average 1.5 meters off the ground, with a horizontal foot print of approximately 2 x 3.5 meters. It is not hard to imagine a system that can provide a rather effective target detection cueing system based on this three-dimensional shape data.

3.4. Functional Classification of Terrain Type

Clearly, target detection is not the only useful purpose of object recognition. If natural and cultural objects can be automatically labeled, then the terrain database can be updated dynamically in a manner similar to the cultural feature overlays provided in ITD data. This data can provide important information for both pre-mission planning and during actual combat operations. Functional areas such as tree canopy, roads, ditches, buildings, rivers, and impassable terrain can be classified, and overlaid on the terrain database map for enhanced visualization. In addition, given one or more target locations (or areas) on the terrain, the terrain can be classified as to whether or not a vehicle (or group of vehicles) is visible from the target location(s). Such information is extremely useful in mission planning and execution, where visibility issues are a tactical concern.

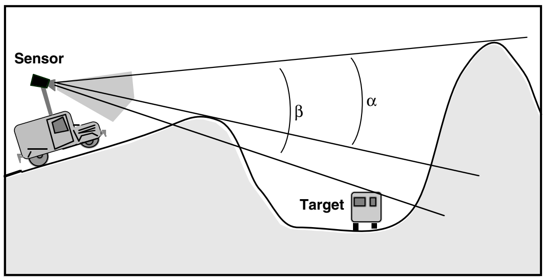

Figure 6. Geometry of 'Hill Cresting'. By servoing on landmark features, the vehicle can position itself so that only the sensor is visible from the target.

Consider the Ft. Hood terrain shown in Fig. 3 where the network of roads through a maze of foliage and vegetation is clearly visible to the human viewer. A very useful enhancement (e.g. for a tactical analysis) to the terrain map would be an automatic analysis of the road network and heavily foliaged areas before the user sees it for the first time.

A long term effort at UMass has been the development of object recognition techniques for both natural objects (Draper, Collins et al. 1989) and cultural objects (Beveridge and Riseman 1995). Using context-sensitive processing strategies (schemas), multiple representations such as 2D lines, points, and regions, as well as 2D and 3D texture and shape measures, can be used to produce accurate functional classifications of terrain maps.

3.5. Fusion of Building Reconstruction and Terrain Modeling.

Umass has been developing site model reconstruction algorithms and IU tools for image analysts as part of the ARPA/ORD RADIUS program. An important function of the site model

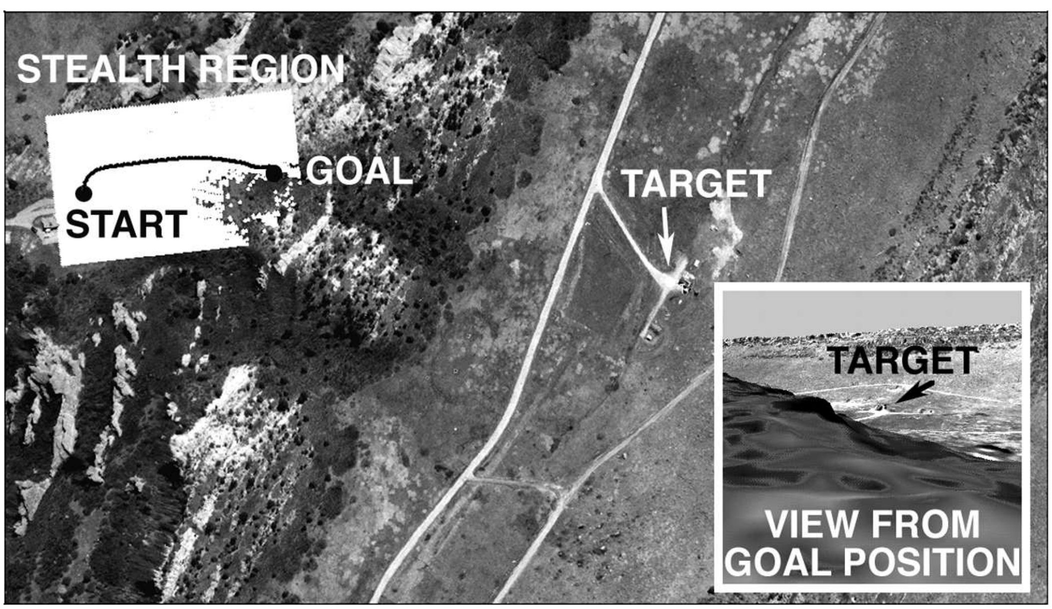

Figure 5. Stealth navigation to a bounding overwatch position for a specific distant target position. Analysis of the terrain allows the vehicle to achieve the goal position without being visible from the target location. The path-planning computation is restricted to the rectangular focus-of-attention region shown in the upper left. White pixels correspond to locations where the vehicle is invisible to observers at the target location.

reconstruction system is to detect closed 2D polygons that correspond to building roofs. These polygons are used to segment the elevation estimates, generated by the terrain reconstruction system, into individual roof facets. Next, a robust plane fitting algorithm is used to compute the orientation and location of each roof polygon; and a set of building "re-assembly" heuristics are used to assemble the building from the fitted planes. Finally, the refined building structures are inserted back into the DEM. More detail on these techniques, and examples of results, can be found in (Collins, Hanson et al. 1996; Jaynes, Stolle et al. 1996) in these proceedings.

4. Fusion of Information into 3D Graphical Visualization System

4.1. Visual Geographic Information System (VGIS)

The technology for producing visualizations of large terrain databases in a fly-through mode has been well-known for some time. However, until fairly recently, the fly-through trajectory had to be determined in advance and it took days to produce the resultant simulation. With new computational engines such as the SGI Onyx processor, this type of visualization can be produced for any viewpoint or trajectory in near real-time (i.e. in seconds). New visualization tools based on this technology are becoming part of the military planning process.

The visualization program being employed in Daedalus is the Virtual Geographical Information System (VGIS) (Koller, Lindstrom et al. 1995). VGIS is an interactive environment for visualizing, navigating, and understanding complex, constantly changing battlefield information databases. VGIS is being jointly developed by the Graphics, Visualization, and Usability (GVU) Center at Georgia Tech and the Battlefield Information Processing Group at the Army Research Labs. The system can be viewed in terms of providing the means for fusing and visualizing two types of data: terrain models comprising elevation data and its corresponding phototexture, along with feature models comprising vehicle CAD models, GIS data (raster layers), and military symbology.

VGIS implements a hierarchical, multi-resolution terrain (elevation) and phototexture database designed to reduce the number of polygons and amount of texture data required during graphical rendering of the visualization (Lindstrom, Koller et al. 1995). Internally, it maintains a quadtree representation coupled with level-of-detail management to cull polygons that are not in the

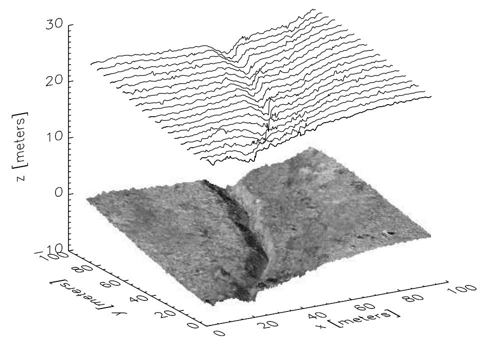

(a)

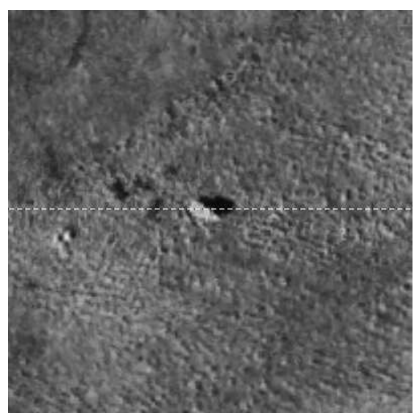

(b)

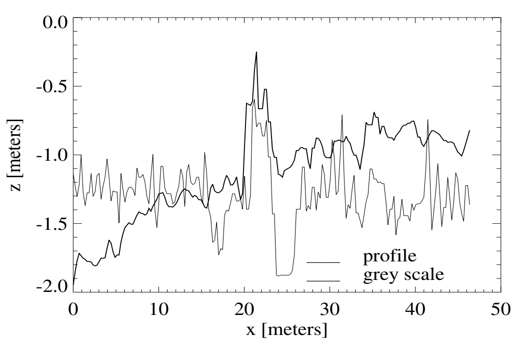

(c)

Figure 7. Obstacle Prediction. (a) Reconstruction (elevation slices and perspective view) of a section of the erosion ditch shown in Fig. 4b. The ditch is approximately 1.5-2 meters wide. (b) Detailed view of the rock visible near the bottom right corner of the perspective view in Fig. 4(b), below the erosion ditch. (c) Elevation and gray scale profile along the dashed line shown in (b); note that the height of the rock, as well as the intensity of the rock and shadow, are clearly consistent in this data.

field of view and select appropriate data resolutions based on terrain quality requirements. Terrain complexity reduction and resolution selection is performed based on several criteria, including distance from viewer, view vector angle, and terrain roughness. For example, terrain regions that are far away from the viewpoint or are relatively flat may be rendered at lower resolutions. Using these techniques, the system is capable of attaining rates of 15 frames per second or more in most cases (on an SGI Onyx processor equipped with a RealityEngine).

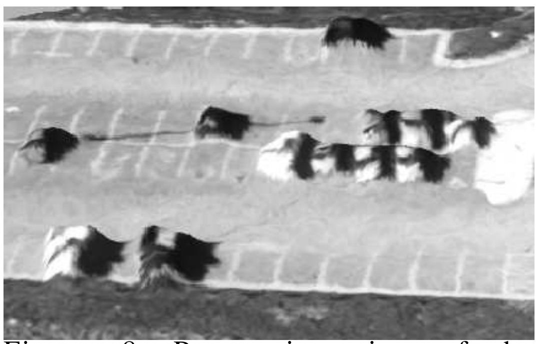

Figure 8. Perspective view of the reconstruction of the Lockheed-Martin parking lot. The parking lot area is shown in the nadir view in Figure 4a. The vehicles in the reconstruction are approximately 1.5 meters high, 3-4 meters long, and 2 meters wide.

The GL and OpenGL graphics programming standards developed by SGI are used as the basis for this visualization capability. VGIS is also built on the Simple Virtual Environment (SVE)(Kessler, Kooper et al. 1995), a toolkit for developing real-time environments for visualizations. Low-level rendering is done via GL or OpenGL while inclusion and management of feature models is done via SVE.

In terms of user interaction, VGIS provides either a workstation window-based interface or a virtual reality interface. Several navigation modes are available, including an aerial map inset facility that allows users to "jump" to a selected location in the terrain and six degree-of-freedom motion (rotation and translation) via a 2D mouse. In addition, the viewpoint may be attached to certain objects such as moving vehicles, allowing the user to “ride” or “fly” through the terrain and cultural features. Simple GIS queries can also be made, for example, to request object names and terrain UTM coordinates. Through the SVE toolkit, vehicle CAD models may be assigned characteristics such as geometry, texture, motion parameters, and additional information (such as RSTA image chips obtained from scout vehicles or other intelligence sources). SVE also provides the mechanisms for displaying and manipulating hierarchical military configurations in the form of standard military symbols. All of these facilities and a variety of other options such as rendering modes and status reports are manageable via pop-up menus.

4.2. Incorporation of Real-Time RSTA Information into Terrain Database.

One of the means of providing an overview of the battlefield is to integrate information from remote real-time sensors mounted on multiple UGV scout vehicles. Since one of the key components of the UGV Demo II program is the intended RSTA functionality (Reconnaissance, Surveillance, and Target Acquisition), it is obviously desirable to immediately incorporate the up-to-date RSTA information into the VGIS terrain database.

Thus, as combat vehicles come into the range of the deployed UGVs and targets are reported, their location, direction, velocity, and IFF status will be overlaid onto the high-resolution terrain model of the battlefield. Vehicles (both enemy and friendly) are added to the visualization and each is rendered in perspective using the appropriate CAD model and placed in the correct position and with the correct orientation. If the vehicles are moving and their velocity is known or can be estimated, that information will be overlaid on the visualization along with its uncertainty. Associated with each identified target is one or more pop-up windows which display the RSTA/ATR image chips received from the scout vehicle; for a concept example, see Figure 9. This capability allows command center personnel to examine the original sensor data as desired.

By fusing all the available information, including both remote sensor and human input, the state of the battlefield can be represented in a dramatic, easy to interpret, 3D representation. The battlefield commander has complete control over the visualization, including the ability to roam at will and to view the situation from any perspective, without being restricted by the original sensor positions. Because the information about vehicle movements is transformed onto the rendered terrain database in an integrated and consistent manner, confusion regarding multiple viewpoints is avoided - despite asynchronous information updates from multiple UGV scout and other ground sensors.

5. Potential Joint UAV-UGV Program

The UAV Tier 2+ and Tier 3- programs (Delaney 1995) have the potential of being extremely important information assets to the UGV program. These programs have been developed as part of a renewed DOD emphasis on battlefield awareness. They are expected to provide virtually continuous aerial reconnaissance of the battlefield by deploying a variety of reconnaissance aircraft, some of which have the ability to stay aloft for over 24 hours and to continuously monitor 40,000 sq. nm per day using visible light and infrared, and eventually synthetic aperture radar sensors. These aircraft are expected to produce datastreams of 50 Mbits per second to satellites and 137 Mbits per second to dedicated ground stations, with a total of over 1 terrabyte of data per UAV per day. Clearly, there is potential for the UAV program to provide aerial image data to the UGV program, allowing construction of terrain maps as little as a few hours before missions.

6. Conclusion

Novel means are required to provide commanders the ability to instantly comprehend dynamic battlefield situations and to react accordingly. The problem of battlefield awareness for the future commander is that he will not be co-located with the multiple, distributed sensors of future systems, and usually will not be intimately familiar with the terrain. The result can be a serious loss in situation awareness. However, in the future, reliable and automatic RSTA systems will be able to detect, identify and track combat vehicles, and automatically add this information into the global visualization.

The Daedalus Battlefield Visualization System will produce real-time 3D graphical visualizations of the evolving battlefield situation. These visualizations will be based on very high resolution terrain maps produced from UAV fly-over images (or other aerial image sources), and will have live situation information from UGVs and other sources of intelligence overlaid on it. By fusing all the available information, including both remote sensor and human input, the state of the battlefield can be represented in a dramatic, easy

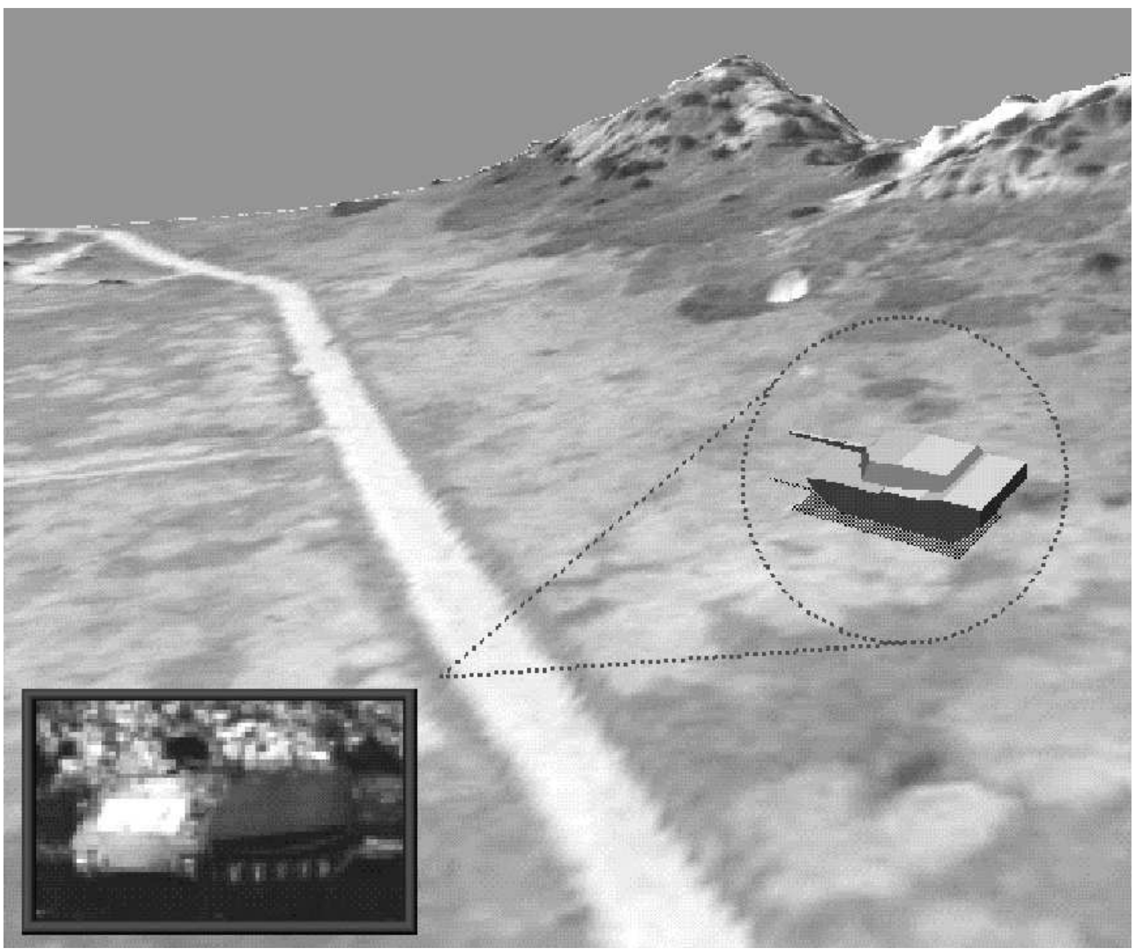

Figure 9. Concept demonstration of incorporating real-time RSTA information into the visualization database. Clicking on the target vehicle brings up one or more windows containing the most recent image chips obtained from a remote sensor. Note that the FLIR image chip shown here does not match the CAD model (there were no APC models in the CAD model database that we have access to).

to interpret, 3D representation. The proposed approach would integrate sensor data from multiple remote sensors into a single, consistent 3D model of the battlefield.

References

Beveridge, J. R. and E. M. Riseman. (1995). “Optimal Geometric Model Matching under Full 3D Perspective,” CVGIP:IU, Vol. 61(3), pp. 351-364.

Collins, R. T., A. R. Hanson, E. M. Riseman, C. O. Jaynes, F. Stolle, X. Wang and Y. Q. Cheng. (1996). “UMass Progress in 3D Building Model Acquisition,” Proc. of ARPA Image Understanding Workshop, Palm Springs, CA.

Delaney, W. P. (1995). “The Changing World, The Changing Nature of Conflicts: A Critical Role for Military Radar,” Proc. of IEEE International Radar Conference, pp. 11-15.

Draper, B. A., R. T. Collins, J. Brolio, A. R. Hanson and E. M. Riseman. (1989). “The Schema System,” IJCV, Vol. 2(3), pp. 209-250.

Jaynes, C., F. Stolle, H. Schultz, R. Collins, A. R. Hanson and E. M. Riseman. (1996). “Three Dimensional Grouping and Information Fusion for Site Modeling from Aerial Images,” Proc. of ARPA Image Understanding Workshop, Palm Springs, CA.

Kessler, D., R. Kooper and J. C. Verlinden. (1995). Simple Virtual Environment (SVE) User’s Guide Version 1.5, Graphics, Visualization and Usability Center, Georgia Institute of Technology.

Koller, D., P. Lindstrom, W. Ribarsky, L. F. Hodges, N. Faust and G. Turner. (1995). Virtual GIS: A Real-Time 3D Geographic Information System, Graphics Visualization and Usability Center, Georgia Institute of Technology, Report# 95-14.

Lindstrom, P., D. Koller, L. F. Hodges, W. Ribarsky, N. Faust and G. Turner. (1995). Level-of-detail Management for Real-Time Rendering of Phototextured Terrain, Graphics Visualization and Usability Center, Georgia Institute of Technology, Report# 95-06.

Ravela, S., B. Draper, J. Lim and R. Weiss. (1995). “Adaptive Tracking and Model Registration Across Distinct Aspects,” Proc. of IROS ‘95, Pittsburgh, PA.

Ravela, S., R. Weiss, B. Draper, B. Pinette, A. R. Hanson and E. M. Riseman. (1994). “Stealth Navigation: Planning and Behaviors,” Proc. of ARPA Image Understanding Workshop, Monterey, CA, pp. 1093-1100.

Ribarsky, W., A. Op den Bosch, J. Bolter and R. van Teylingen. (1994). “Information Visualization and Analysis using Virtual Reality,” IEEE T-CG&A, Vol. 10.

Schultz, H. (1994). “Terrain Reconstruction from Oblique Views,” Proc. of ARPA Image Understanding Workshop, Monterey, CA, pp. 1001-1008.

Schultz, H. (1995). “Terrain Reconstruction from Widely Separated Images,” Proc. of SPIE, Orlando, FL, pp. 113-123.