Target Identification Using Geometric Hashing and FLIR/LADAR Fusion[1]

Alexander Akerman III and Ronald Patton
Nichols Research Corporation
3452 Lake Lynda Drive, Suite 190
Orlando, FL 32817
akermana@puzzler.nichols.com

Walter Delashmit
Lockheed Martin Vought Systems
P.O. Box 650003
Dallas, TX 75265-0003
delashmi@lvs-emh.lvs.loral.com

Dr. Robert Hummel
New York University
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, NY 10012
hummel@relaxation.cs.nyu.edu

Abstract

The primary objective of this project is to provide the Surrogate SemiAutonomous Vehicle (SSV) with a Demonstration II capability of performing automatic target recognition/identification (ATR/I). Detected objects of interest will be imaged with FLIR and/or LADAR sensors so the ATR/I algorithms must be compatible with either sensor, as well as exploit the synergy of processing both sensors simultaneously. Our approach does not rely upon the precise coregistration of multiple sensors, but rather performs geometric hashing on the individual FLIR and LADAR images.

The geometric hashing software, originally developed for 2D SAR and FLIR imagery, has been extended to also accommodate 3D Ladar range and intensity imagery. The 2D hashing software was modified to allow up to ten dimensions. Currently, a 4D scheme is being used which represents (x,y) position, range, and intensity features.

A Laplacian pyramid algorithm has been developed to fuse Ladar range and intensity imagery. Although previously used with dissimilar sensors (e.g., FLIR and TV), the algorithm proves maximally efficient with pixel coregistered imagery, as in the case of the two Ladar modes. Both target and background contrast and internal structure are significantly enhanced, thus facilitating target segmentation and feature extraction. The algorithm can be implemented with the existing pyramidal processors provided by Sarnoff Labs for RSTA image stabilization.

The FORTRAN/VAX software has been rewritten into C/UNIX and installed on a SUN SPARC processor within the Surrogate SemiAutonomous Vehicle (SSV). The C/UNIX software provides at least a 30-fold decrease in execution time and a 20-fold increase in model storage capacity.

Operational target recognition was performed during Demo C in July 1995 using both actual FLIR and synthetic LADAR imagery. 100% correct classification was obtained on the six target, 12 pose synthetic LADAR imagery which we had generated using the LARRA/SAIL and BRLCAD models. A subsequent laboratory experiment using the Demo C 66 FLIR target model set achieved an 86% correct classification of 56 unknown targets.

---

[1] This work was funded by the ARPA/ISO Image Understanding program and was performed under ARO contract DAAH04-93-C-0049.

1. Introduction

Geometric hashing is the fundamental technique which we have applied to the military ATR problem [Akerman, et al., 1992]. Conceived by researchers at the NYU Courant Institute [Lamdan and Wolfson, 1988], hashing represents an object by a collection of points, which are then matched to similarly constructed models. The matching is accomplished by iteratively selecting pairs of points, placing them in a Euclidean geometry coordinate system, concurrently translating and rotating all other object points to the same geometry, and then counting the number of occurrences of object and model points in the same cell.

Geometric hashing is particularly appealing since it can be very efficiently implemented with parallel processing. An unknown object can be simultaneously tested against thousands of models, including specific orientations/states of each target [Bourdon and Medioni, 1990].

Our contribution to hashing algorithms has been the application of the technique to military targets in Synthetic Aperture Radar (SAR), Forward Looking Infrared (FLIR) imagery, and LADAR imagery. During 1988-1990, we developed a SAR point extractor, determined thresholds and tolerances for SAR point matching, and created software for simultaneous multiple model testing [Akerman and Patton, 1990].

In 1991, we began an investigation of FLIR imagery hashing, with particular emphasis on the extraction of robust hash points from the targets. Algorithms were developed to select points that represented the target's geometrical structure and that were thus stable and repeatable under various radiometric conditions, due both to the external environment and to the target itself. These algorithms entailed first extracting significant contours corresponding to the target's key components (tread, turret, etc.). The hash points are obtained from the end points, intersections, and key curvature of those lines.

In 1993, we refined the FLIR-associated hashing algorithms under a contract with the Army Night Vision and Electronic Systems Directorate [Akerman et. al, 1993]. Specifically, the algorithms were extended to second generation FLIR imagery which provides much greater detail of target internal structure.

As shown by Figure 1, our algorithm architectures will fuse the results of the FLIR and LADAR hashing by using a Piecewise Level Fusion Classifier (PLFC). Such fusion is based upon the work of Thomopoulos, [1987], which we are currently implementing on another project for the fusion of FLIR and MMW Radar data.[2]

Target boundaries are extracted as an intermediate step in the determination of hash points (per edge curvature and edge intersection). Hence, we also perform multisensor fusion using combined FLIR and LADAR boundaries to perform Recognition-By-Components (RBC). We extend the work of Lowe [1985], Biederman [1987], and others to provide viewpoint invariant recognition using perceptually organized features of geometric components. A Bayesian reasoning structure is used to fuse the results from the Hashing PLFC and the RBC algorithms.

As indicated, work thus far completed is shown by the solid blocks in Figure 1. As the prime contractor, Nichols Research Corporation (NRC) has developed the overall architecture and the front end FLIR and LADAR image enhancement algorithms, as well as the geometric hashing codes. Our techniques for LADAR image enhancement fuse the pixel coregistered range and intensity images by merging the individual levels of a Laplacian Pyramid decomposition of each of the LADAR images. The FLIR image enhancement uses classical histogram equalization nonlinear mapping, and gradient sharpening techniques.

---

[2] This work is being performed for the US Army CECOM Night Vision and Electronic Systems Directorate under contract DAAB07-94-C-M034 as part of the Target Acquisition/Target Recognition Phase II SBIR A92-042. POC is Dr. Tom Witten at (703) 704-3052, witten@nvl.army.mil

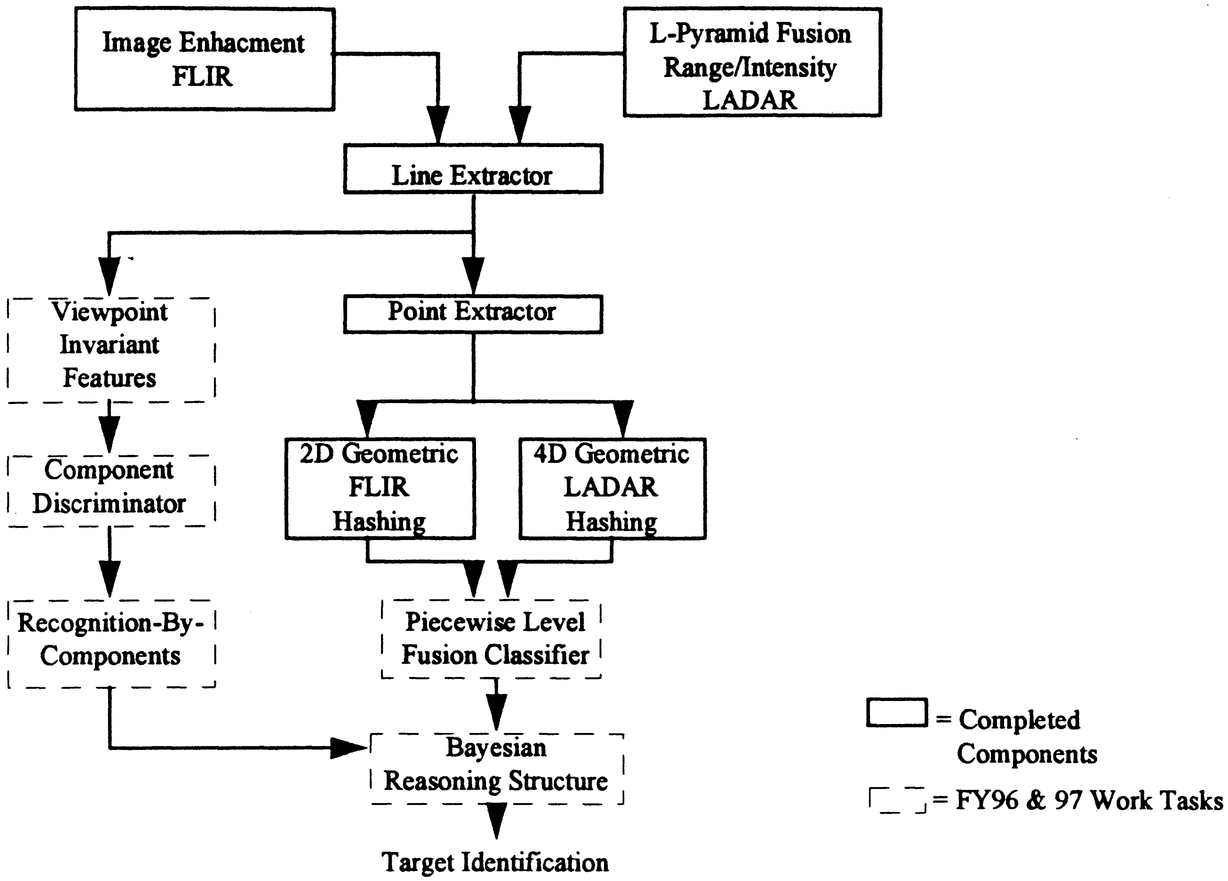

Figure 1. Overall Architecture for Multisensor Fusion Using FLIR and LADAR Identification

2. Ladar Intensity and Range Image Fusion

We have previously developed an algorithm for fusing FLIR and TV, (as well as laser intensity and TV) images into a single fused image [Akerman, 1992]. In essence, the algorithm represents each image as a Laplacian pyramid [Burt and Adelson, 1983], and then combines the individual sensor representations one level at a time using an appropriate pixel select criteria [Toet et al, 1989].

The resultant image quality is significantly dependent upon pixel coregistration between the two individual images. Hence, one might expect optimal results in the case of Ladar range and intensity imagery from the same sensor and thus exactly pixel coregistered. Figures 2 and 3 indeed illustrate that such an image fusion does provide a significant enhancement in target and other scene detail.

The upper subimage is the Ladar intensity return. Notwithstanding numerous pixel "dropouts," this image provides good internal detail of objects. However, those objects often lack distinct borders and instead blend with their immediate background. In particular, note the wheeled object in the lower left quadrant of Figure 2. It could be very difficult for an automatic target recognizer to discern that it is a truck.

The middle subimage is a transform of the Ladar range image, in which a zero gray scale (Black) is the ground plane and higher Gray Scale values represent increasing height above the ground plane. For this representation, internal object structure is

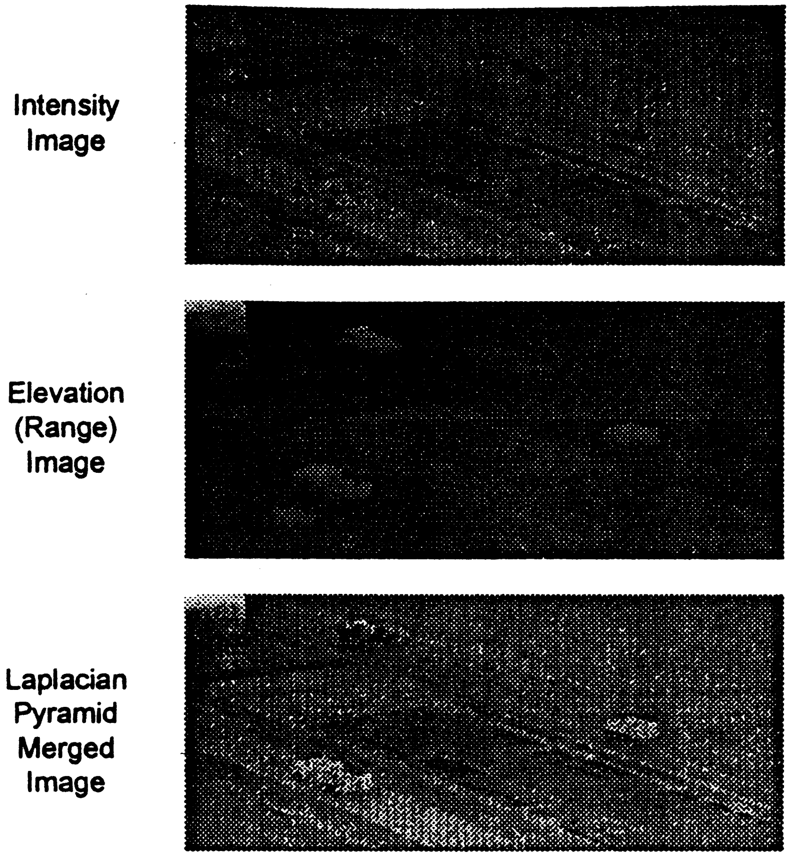

Intensity Image

Elevation (Range) Image

Laplacian Pyramid Merged Image

Figure 2. Ladar Image Enhancement By Image Fusion

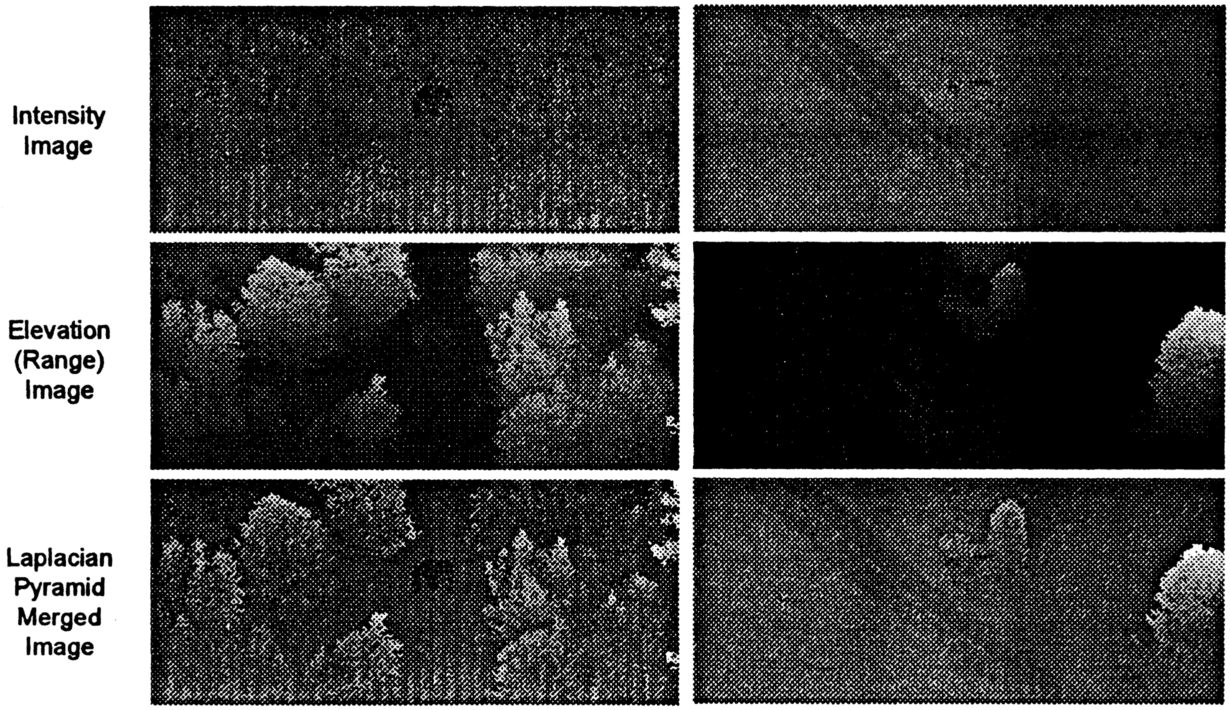

Intensity Image

Elevation (Range) Image

Laplacian Pyramid Merged Image

Figure 3. Additional Examples of Ladar Image Enhancement

minimized, particularly near the ground plane. However, the overall shape silhouette is enhanced. The bottom subimage portrays the resultant image fusion. The wheeled object in Figure 2 is now clearly a truck. The other vehicles in that image are also more distinct. Figure 3 presents two other scenes. For the one on the left, the target was already very distinctive in the intensity image, so the merging provides no significant enhancement. In comparison, however, note the detail in the foliage of the merged image as compared to that in either of the intensity or elevation images.

For the scene on the right side of Figure 3, note the truck next to the tree. The truck is not very distinct in either the intensity or the elevation images, but is "pulled out of the mud" in the merged image.

In all instances that we have thus far processed, the merged image never has an object of lower image quality than that of the intensity or elevation image alone. Often, there is a very significant improvement as we have shown. Hence, all of our geometric hashing algorithms are being applied only to Ladar imagery that has merged both the intensity and range signatures.

Figure 4 illustrates the application of the Rule-Based Line and Point Extraction algorithms applied to Ladar imagery of four target types. This imagery was collected with a Loral Vought diode umped laser radar operating at 1.06μ, with a 0.4mr horizontal and vertical angular resolution, and with a 0.15m range resolution. The targets are at ranges of 300-400m and at depression angles of 14-18°.

The line extraction algorithms were applied to both the Laplacian Pyramid fused image (of the Ladar intensity and elevation images) and to the elevation image alone. Note that while the Laplacian fused image yields significant internal detail, it does not capture all of the target's exterior boundary. (Although not shown, this deficiency is significantly worse when only the Ladar intensity image is used). Conversely, the Ladar elevation image yields a good segmentation of the overall target shape but loses much of the internal target detail.

When both line extractions are combined together, all of the key geometrical components of the target are distinctly outlined. Note also that there are no extraneous lines on the target, except when there are obscuring clutter artifacts. Hence, the line segmentation provides a very robust geometry for the extraction of the hash points, which are also shown in Figure 4.

3. LADAR Hashing Models and Matching Results

3.1 LARRA/SAIL Synthetic LADAR Imagery

Synthetic LADAR range images were generated by NYU from the Ballistic Research Laboratory (BRL) computer aided design (CAD) models for six tactical vehicles (targets) using the Laser Radar Recognition Algorithm (LARRA)/Synthetic Assembly Image Layout (SAIL) modeling code. The targets selected were the M60A3 tank, M113 armored personnel carrier (APC), M35 truck with a rack or a canvas cover, M35 truck without a rack or canvas cover, HMMWV troop version with the conventional sloped rear and HMMWV cargo version which has a square back. The rationale used for selection was the availability of the appropriate BRLCAD models for unrestricted access. These images were generated at a sensor depression angle of zero degrees to correspond to the Unmanned Ground Vehicle (UGV) scenario. Images were created every 15 degrees from zero to 360 degrees. The original images were created using a resolution of 0.05 milliradians (mr) in both azimuth and elevation with each pixel corresponding to a ray trace. All output images are in the Khoros viff format. For these six targets, images at every 30° were selected to use to build up a data base of 72 models consisting of 12 orientations for the 6 targets.

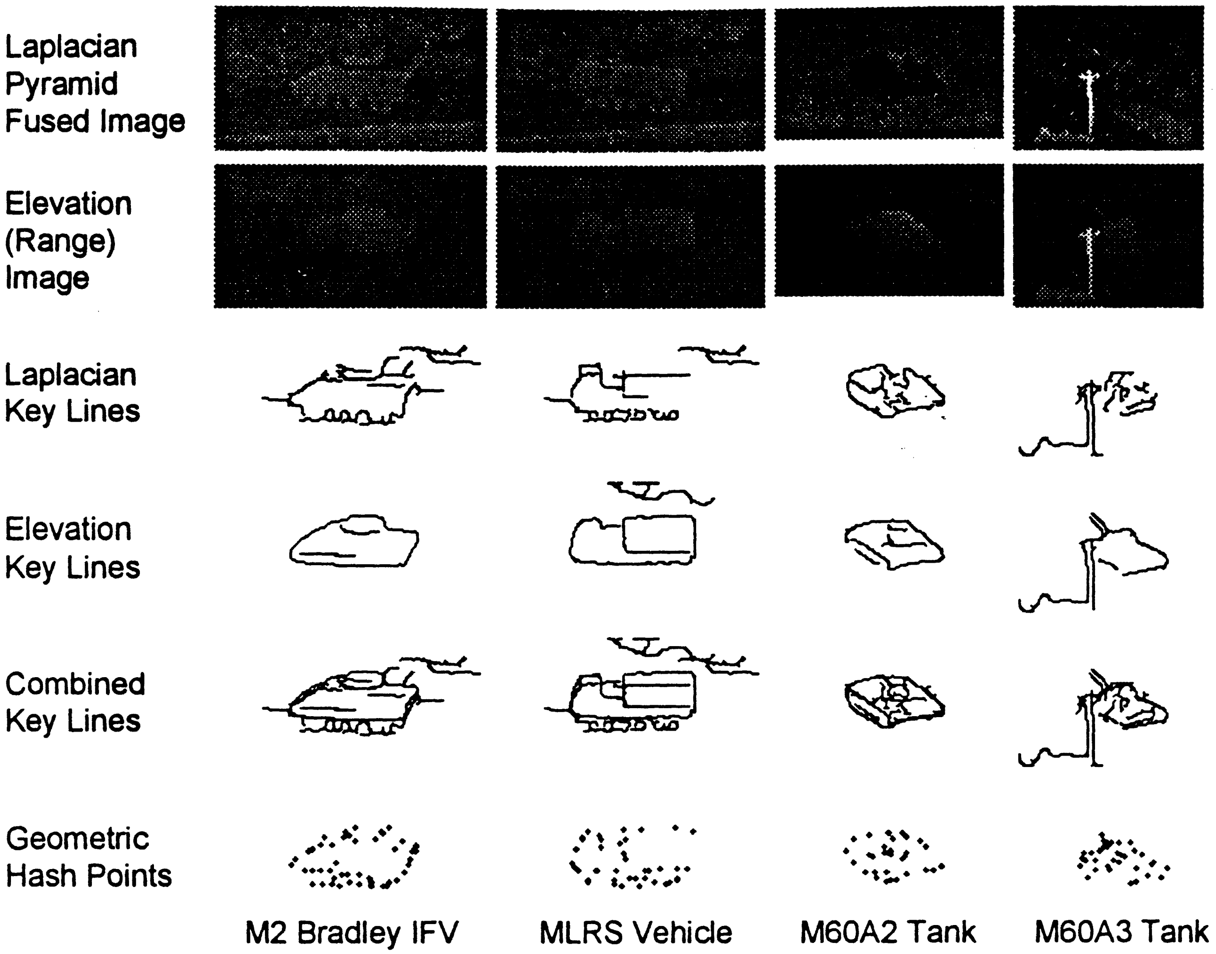

Figure 4. Baseline Feature Algorithms Applied to Ladar Imagery

3.2 Hash Point Models

Image chips, line extraction/segmentation and extracted hash points are shown in Figure 5 for three different orientations of the M113 APC corresponding to aspect angles of 285°, 15°, and 270°. The white pixels in the image chips correspond to range data drop outs where a very large nontarget value was recorded. It should be noted that the line and point segmentation works even for cases where a complete target outline in terms of edge/line structure is not obtained. It can also be clearly seen in this figure that line end points, line intersections and points of curvature have all been extracted for the models.

Similar results are also shown in Figure 6 for three different orientations of the M60 tank corresponding to aspect angles of 150°, 90°, and 75°. For this case, the tank outline is less complete than for the M113 due to much missing structure along the bottom of the tank tread. In addition, one background point was picked up for the M60 tank at 90° aspect, but this single extraneous point will have little or no effect on match results.

Visual comparison of Figures 5 and 6 clearly shows dissimilarities in the extracted point structure for the two different targets (M113 and M60). This dissimilarity in the extracted features, i. e., hash points, is the basis on which the geometric hashing algorithms are used for target identification.

3.3 Test Results

The geometric hashing algorithms were tested in a real-time operational scenario during Demo C on July 27, 1995. Due to operational time constraints for the entire demonstration of our work and the work of other contractors, only five LADAR images were processed. The five images corresponded to a M60 tank at 270° aspect, a M113 APC at 15° aspect, a HMMWV cargo version at 15° aspect, a M35 truck with the rack and canvas top at 300° aspect, and another M60 tank but with this tank at 120° aspect. The results of this real-time operational test was that the geometric hashing algorithms produced 100% correct target recognition. It should be noted that this realtime testing was a recognition task rather than an identification task since the available target set did not support identification, i. e., there were not multiple types of a single class of targets (e.g., M60, M1, T62, T72, tanks, etc.).

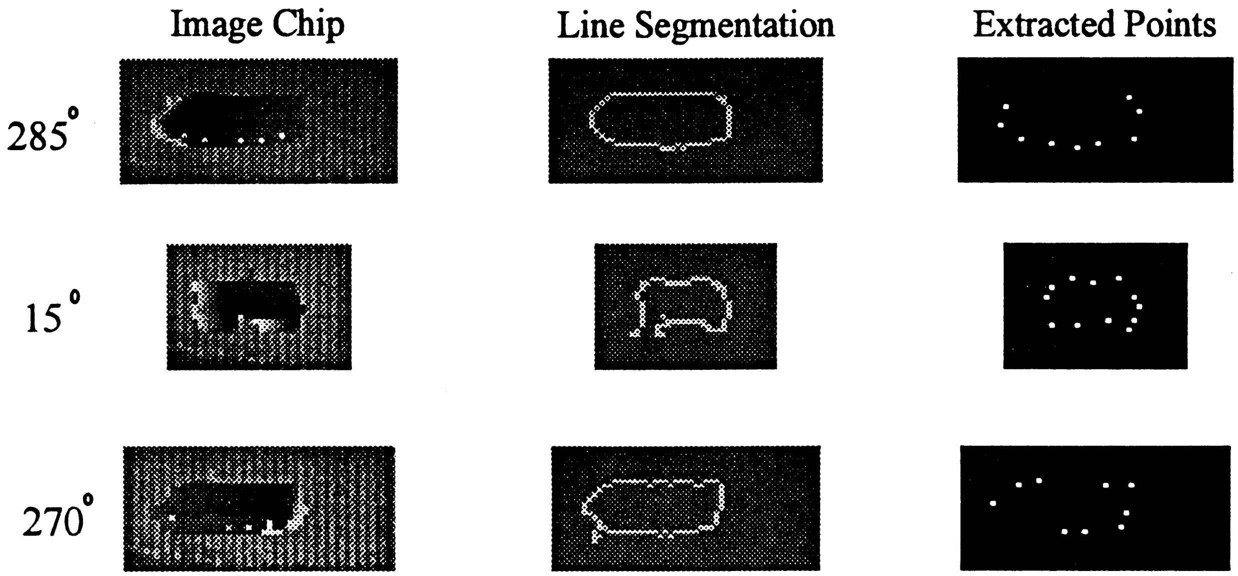

Figure 5. LADAR Line Segments and Extracted Hash Points for M113 APC's

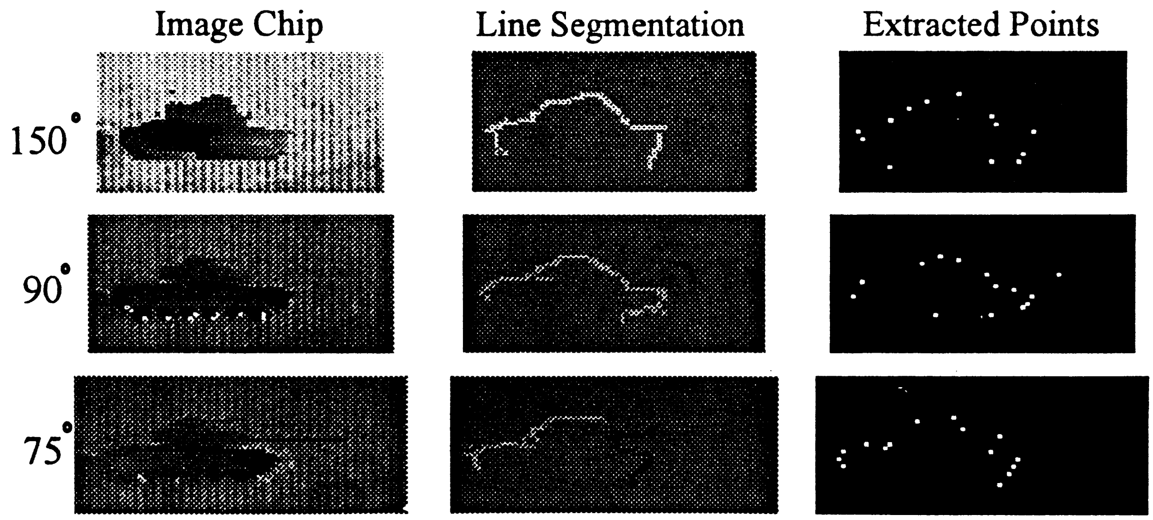

Figure 6. LADAR Line Segments and Extracted Hash Points for M60 Tanks

4. SSV Sun SPARC Software

Existing FORTRAN geometric hashing code was converted to C and enhanced to run in real time on the Sun SPARC computers which are available on the SSVs. This C UNIX software was delivered to Lockheed-Martin and support was also provided to install and checkout the software to ensure proper operation of the code. This effort was crucial to the successful demonstration for Demo C as discussed above.

4.1 Code Conversion and Enhancement

The original FORTRAN code was converted to C to run under an UNIX Sun workstation. A LINUX version (PC UNIX) is also available. Various enhancements and updates to the code were made during the conversion process. Some of these changes were made to make the code more efficient for real time implementation, some were made as a result of SSV requirements specified in the Lockheed-Martin Software Configuration Control Document [Severson, 1995], some were made to make the code more portable and essentially machine independent, some were made to increase the numerical precision of the code, and some were made to update the original version of the code. This code development refers not only to the code hosted on the SSV for real time usage, but also to code which was developed to produce the hashing model files. Structured C programming techniques were used with extensive documentation embedded in the actual code.

To enhance real time operation, the model file is read in from memory at the start of the program by the executive controller and maintained in memory during system operation. This results in significant speed enhancement. This software is configured so that the model file can be stored and read as either a binary file or an ASCII file. It is preferable to use binary files which results in reduced storage requirements and enhanced speed.

4.2 Real Time Architecture

The real time architecture was designed to be compatible with the Lockheed-Martin Software Configuration Control Document. As such all input and output is performed in the main program body and the geometric hashing code is called as a function from this main program. This main program provides the hashing points from the feature extraction software which have been hosted on the SSV Sun SPARC computer. The model file is also read in by this main program. The output results are transferred back to the main program which then outputs the classification results for target class and confidence measure and the image chip along with the line segmentation and point images for display purposes. No calls to "exit" are made internally to any of the hashing routines.

4.3 Efficient Hash Table Generation

The hash table and model file generation is all performed internally to memory for all operations. No specialized commercial software is used for any of these functions. A previous version of the code used the CINDEX software developed by Trio Systems Inc. This was required in the previous system based on the hosting processor. However, this specialized software is somewhat machine dependent and would have to be licensed for any machine on which the software were to be hosted. These complications were avoided by using a structured key model technique with all computations being performed internally.

The previous FORTRAN keyed-access file used to store the patterns associated with each master-slave point pair was composed of an (x,y) coordinate that uniquely specified a record in the file. Interactive disk accesses are not appropriate for real time operation and would be inconsistent with the software configuration requirements. Thus this process was replaced with a memory resident equivalent while maintaining a similar interface for access to individual records.

5. FLIR Hashing Models and Matching Results

For Demo C conducted at Lockheed-Martin, Denver, in July 1995, the NRC/LV team provided a suite of hashing-associated codes for the SSV SPARC processors. Two separate sensor hash tables were also created: (1) a 72 model LADAR set, which was described in Section 3, and (2) a 66 model FLIR set which is now addressed.

5.1 FLIR Imagery Used for Model Building

Unlike the LADAR models which were derived from synthetic imagery, the FLIR models were produced from imagery collected by Lockheed Martin on 6 and 7 October 1994 [Munkeby, 1995]. The 3-5 micron Amber FLIR was the same type of sensor as that used on the operational SSV's. This FLIR has a square 2.73° x 2.73° narrow FOV resolved into a 256 x 256 pixel array, and thus a resolution of 0.19 milliradians.

The October 6 and 7 data collections were specifically planned to generate training data for model building, with all targets at a fixed 961 meter range. All targets were on flat, level ground, and the targets were rotated only in azimuth in 30° incremental steps. Each scenario consisted of three targets at a fixed azimuth orientation with one wide FOV image taken of all three targets simultaneously and ten each sequential narrow FOV images for each of the three targets. (For building the FLIR model hash table, we used only one narrow FOV image for each target).

The first three collected targets consisted of the M113 APC, the M35 truck, and a HMMWV. Ten scenarios (500-509) were collected on 6 October 1994 and four more scenarios (510-513) on the following day. The second target set was collected entirely on 7 October for ten scenarios (520-529). It consisted of an M543 Wrecker, an M60 Tank, and another HMMWV. The data collection had to be terminated prematurely due to rain.

The entire data collection occurred under marginally-acceptable weather conditions (cold, cloudy, and increasingly overcast) with deteriorating weather (impending rain) washing out much of the internal target detail. Such environmental conditions do not correspond with those typical in July, which was when Demo C was performed.

5.2 Hash Point Models

Figures 7 and 8 provide the twelve views and corresponding point models for two of the targets: the M113 APC and the M60 tank (for the latter, the 330° and 300° orientations are absent, because that data was not collected).

Inspection of Figures 7 and 8 shows that in many instances the point models do not exactly mimic the target geometry. This is due mainly to the lack of sufficient quality in many of the images. Also, the Line and Point Extractors are not perfect, even when the image quality is very good. Notwithstanding these degraded point representations, they are nonetheless sufficient in almost all instances to provide a unique representation by target type and orientation. Hence, it should not be too surprising that excellent classification results were obtained against this 66 model set, as discussed next.

5.3 FLIR Hashing Test Results

The 66 FLIR model hash table was initially tested at Demo C, for which two of the three "unknown" targets (M113 APC, HMMWV, and M35 Truck) were correctly recognized. The very limited Demo C schedule did not allow additional target types to be tested. The match criteria were:

100 percent model points used
1 pixel mismatch tolerance
100 percent live points used
1.4 average pixel mismatch
50 percent live points matched
10° in-plane rotational angle disparity

To quantify more thoroughly the FLIR hashing performance, an extensive laboratory experiment was subsequently conducted using the same 66 model hash table. The original images from which those models were derived were input into the overall hashing suite (except the second HMMWV target was not input). Hence, there were 56 trials, for each of which there was a complete process of image enhancement (but without spatial sharpening), line extraction, and point extraction. Those automatically extracted points (without any human adjustment) were then tested against the 66 model hash table. The same matching parameters as those of Demo C were used, except that the percent of live points to be matched was reduced to 20%. Tables 1 and 2 give the resultant classification matrices, both for absolute numbers and corresponding classification probabilities. The overall average classification probability is 86%, which includes all trucks and wreckers being correctly classified.

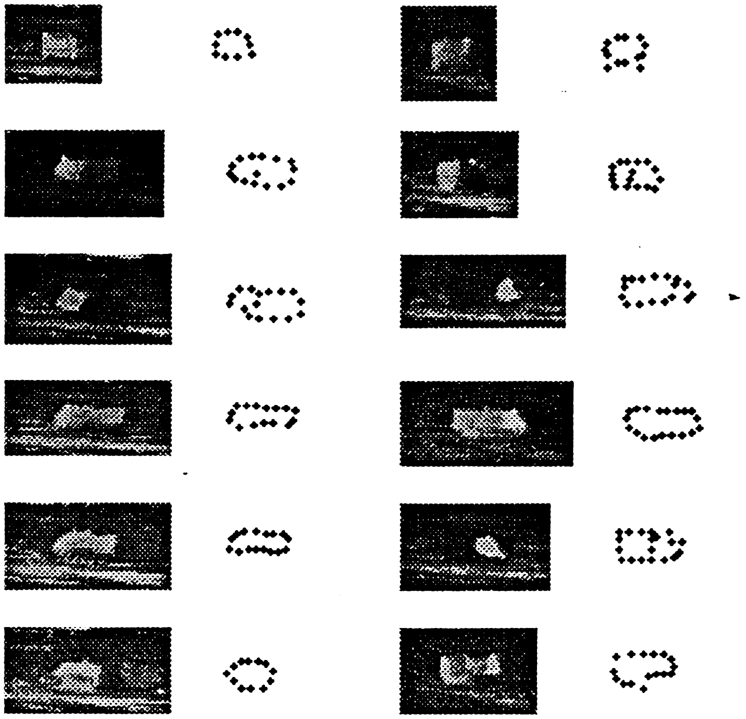

Figure 7. M113 APC Images and Point Models

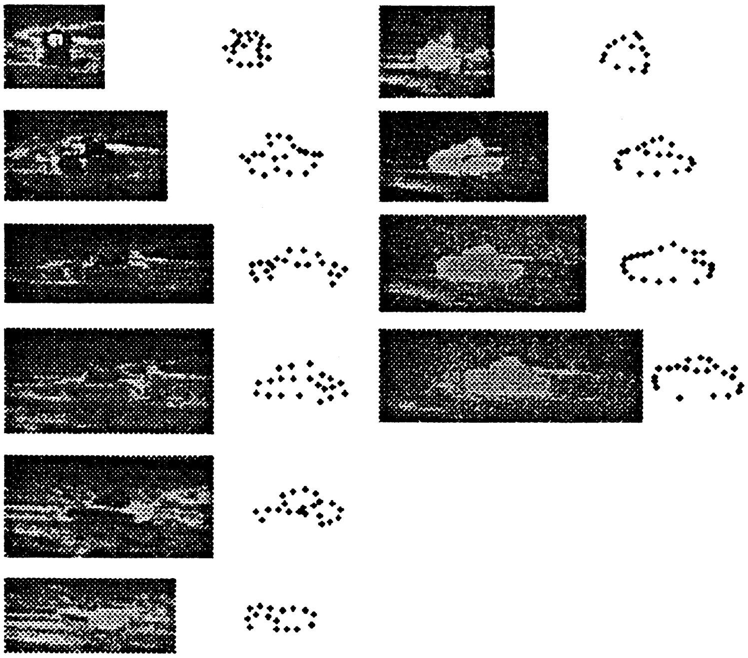

Figure 8. M60 Tank Images and Point Models

Table 1. Overall Classification Matrix

	APC	Truck	Tank	Wrecker	HMMWV
APC (M113)	9	3	----	----	----
Truck (M35)	----	12	----	----	----
Tank (M60)	----	1	9	----	----
Wrecker (M543)	----	----	----	10	----
HMMWV	1	----	1	2	8

Table 2. Overall Classification Probability Matrix (P_{cc} = 86%)

	APC	Truck	Tank	Wrecker	HMMWV
APC (M113)	0.75	0.25	0	0	0
Truck (M35)	0	1.00	0	0	0
Tank (M60)	0	0.10	0.90	0	0
Wrecker (M543)	0	0	0	1.00	0
HMMWV	0.08	0	0.08	0.17	0.67

6. Conclusions, Observations, and Recommendations

The UGV RSTA program provided the first opportunity to quantify geometric hashing performance in a military context, both with respect to various target types and on an operational platform, the SSV. Although several laboratory and field experiments have now been conducted, no overall assessment should be made until the results of Demo II and the NVESD and independent contractor laboratory evaluations can be reviewed. To date, the geometric hashing algorithm has exhibited very favorable performance and was the only recognition/identification software incorporated into the Demo II SSV's.

The 3-5μ FLIR produced imagery much different, and often of much lower quality, than that of 8-14μ FLIR's typically used for ATR. Hence, the data collections were insufficient for suitable model building. Too much emphasis was given to difficult conditions (e.g., vehicles on hills, obscuration, etc.) at the expense of not first generating a comprehensive target data base at precise orientations, ranges, and for a variety of times of day, year, and illumination conditions. The associated data basing was marginal, due to the ambitious data collection objectives and very limited resources made available to meet those objectives.

A shortcoming of the program was in not providing target imagery, either synthetic or real, to demonstrate a major RSTA requirement of target identification. That is, there were no foreign targets, even though operational scout personnel consistently stressed the need for a UGV to discriminate friendly from enemy vehicles.

The uncertainty of whether there would be a Ladar sensor onboard the UGV, much less what type of Ladar it would be, caused many discontinuities in developing the ATR algorithms, particularly those that involved the fusion of Ladar with other sensors.

In the case of the geometric hashing algorithms, outstanding Ladar ATR performance was achieved against the synthetic LASER+ and LARRA/SAIL model-generated imagery. It is unfortunate that these models could not be tested operationally in the same manner as the FLIR ATR, i.e., on the Demo II SSV's. (A self evaluation of the Ladar ATR is planned using the MICOM/Eglin Ladar tower test imagery, which Loral Vought collected previously as part of the LOCAAS program. Hence, some modicum of an assessment against actual imagery may yet occur).

The most important contribution of our team to the UGV RSTA program has been to transition the hashing software from a non-realtime, laboratory code to near real-time software resident on a SPARC workstation and thus operable in a military vehicle like the SSV.

If our project had been funded as originally intended (three equal amounts in each of three successive years), then additional algorithm development in the fusion domains could have been achieved by now. Such developments would have included operational SSV software of the same form as the geometric hashing software. Instead, much of that fusion work has had to be deferred until after the completion of Demo II.

The prime contractor, Lockheed-Martin performed admirably considering the myriad of technologies it had to integrate (into the SSV) from the many organizations, (university, government, and contractor). In the case of RSTA, Lockheed-Martin suffered a very large personnel turnover and yet managed to assimilate the geometric hashing algorithms and software into the RSTA suite with only modest support from us. Although their relationship with all the RSTA co-contractors was very good, it sometimes became somewhat unclear as Lockheed-Martin switched back and forth from being just the system integrator to also being one of the RSTA algorithm developers.

Notwithstanding any of the difficulties cited above, this ARPA initiative was an exciting and enjoyable experience that significantly pushed forward the ATR state of the art. We were very pleased to be a part of it.

References

Akerman, A., Patton, R., Smith, S., and Williams, A., FLIR and MMW Algorithms to Detect and Classify Stationary Targets, I-MATH Associates, Inc. Report 93IR36, Contract DAAB07-93-U-0006, USA NVESD, Ft. Belvoir, VA, 8 November 1993; also in Proceedings of the 4th ATR Systems and Technology Symposia, ATRWG PR-94-001, Vol. I, pp. 61-90, November 1994.

Akerman, A., "Pyramidal Techniques for Multisensor Fusion," SPIE Proceedings 1828-15, O/E Tech 92 Symposia, November 1992.

Akerman, A., Patton, R., and Mayes, W., Parallel Processing Application for Avionics, USAF-WL-TR-92-1063, May 1992.

Akerman, A. and Patton, R. "Model-Based Vision Using Geometric Hashing," 1990 Applied Imagery and Pattern Recognition Workshop, SPIE Proceedings 1406-07, November 1990.

Biederman, I., "Recognition-By-Components: A Theory of Human Image Understanding," Psychological Review, Vol. 94, No. 2, pp. 115-147, 1987.

Bourdon, O. and Medioni, "Object Recognition Using Geometric Hashing on the Connection Machine," Proceedings of the ICPR, 1990.

Burt, P.J. and Adelson, E.H., "The Laplacian Pyramid as a Compact Image Code," IEEE Trans COM, 31, (4), April 1983.

Lamdan, Y. And Wolfson, J.J., "Geometric Hashing: A General and Efficient Model-Based Recognition Scheme," Proceedings of the Second International Conference on Computer Vision, pp. 238-249, December 1988.

Lowe, D.G., Perceptual Organization and Visual Recognition, Kluwer Academic Publishers, ISNB 0-89838-172-X, 1985.

Munkeby, S.H., RSTA September 94 Data Collection, SSV Report, ARPA Order 8353, Martin Marietta Astronautics, January 31, 1995.

Severson, W., "RSTA Interface Files," Martin Marietta Technologies, Inc., 27 February 1995.

Thomopoulos, et al, "Optimal Decision Fusion in Multiple Sensor Systems," IEEE Trans AES-23, No. 5, pp. 644-653, September 1987.

Toet, A., Van Ruyven, L.J., and Valeton, J.M., "Merging Thermal and Visual Images by a Contrast Pyramid," Optical Engineering, 28, (7), July 1989.