REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED *(From - To)*
30-11-2009	Final Program Report	01-Jul-2007 – 30-Nov-2009

4. TITLE AND SUBTITLE

Fusion Metrics Final Report

5a. CONTRACT NUMBER
FA9550-07-C-0065

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Steven Haron, David Fay, Paul Ilardi

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

BAE Systems Advanced Information Technologies   6 New England Executive Park Burlington, MA 01803

8. PERFORMING ORGANIZATION REPORT NUMBER

TR-2533

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR
875 N. Randolph St. Rm 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution Statement A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

In the Fusion Metrics program, a multi-sensor testbed was developed and an associated multi-sensor fusion system was demonstrated in an initial data collection. In developing this testbed, algorithms developed in part in previous Air Force research programs were developed into a single platform. A demonstration included moving vehicle targets captured in multiple image sensors, and the real-time fusion of imagery to detect and maintain track on targets. Next, the same algorithms were developed into a modular software fusion testbed to demonstrate rapid prototyping of fusion methodologies via this fusion metric software testbed. Multi-sensor video was fused and fusion classifier algorithms were evaluated via statistical metrics. The software fusion metrics testbed represents a way to rapidly assess new fusion algorithms, and to perform parameter tuning to optimize fusion systems.

15. SUBJECT TERMS
Multi-sensor Fusion, Fusion Metrics, Fuzzy ARTMAP, Video Fusion

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Steven Haron
a. REPORT UNCLASSIFIED	**b. ABSTRACT** UNCLASSIFIED	**c. THIS PAGE** UNCLASSIFIED	UL	13	**19b. TELEPHONE NUMBER** *(include area code)* 781.262.4588

Fusion Metrics Final Report

TR-2533

Contract Number:

FA9550-07-C-0065

Reporting Date:

November 30, 2009

Prepared For:

Air Force Office of Scientific Research

Prepared By:

Steven Haron

steven.haron@baesystems.com

David Fay

david.fay@baesystems.com

Paul Ilardi

paul.ilardi@baesystems.com

BAE Systems

Advanced Information Technologies

6 New England Executive Park

Burlington, MA 01803

(781) 273-3388

TABLE OF CONTENTS

1 INTRODUCTION

The Fusion Metrics program was an AFOSR research effort at BAE Systems Advanced Information Technologies (AIT) spanning from July 2007 to November 2009. Original goals focused on developing a fused video quality metric based on multi-target performance tracking. As part of this research, BAE assembled a multi-sensor suite, comprising a Hitachi RGB camera and a Hitachi low-light EMCCD camera (both with Fujinon zoom lenses), a Goodrich/Sensors Unlimited short-wave IR (SWIR) camera, a HEAT uncooled long-wave IR (LWIR) camera, and a Newcon Optic laser range finder (LRF), integrated on a Directed Perception PTU-D300 rugged pan-tilt unit using a mechanical mount and serial-to-zoom lens controller. Real-time video image fusion, target detection, multi-target tracking, and sensor resource software modules were integrated with the sensor suite, and were laboratory tested. In addition, a daytime tower test was conducted using a simpler dual-sensor suite. In the tower test, two moving (crossing) vehicles were tracked in the vicinity of a stationary shed. During this test, sensor data was collected and real-time fusion and tracking processing was demonstrated.

In 2009 the focus of the Fusion Metrics program was adjusted to develop a fusion metrics software platform for metric assessment that would leverage a data set associated with an Air Force Research Laboratory (AFRL) program. Using this dataset, BAE applied fusion algorithms in Monte Carlo tests to statistically evaluate utility of fusing information in the data. A key outcome was a modular fusion software testbed that showed utility for rapid prototyping of fusion methodologies and to produce metrics for fusion analysis. BAE presented an overview of this fusion testbed at the Annual PI Review in October, 2009.

2 DEVELOPMENT OF HARDWARE PLATFORM FOR FUSION ANALYSIS

An initial focus in the Fusion Metrics program was to produce a data set that could be used to fuse data from a multi-sensor test suite and then compare automated fusion performance to that of a human operator instructed to manually fuse data. In the process of analyzing automated fusion relative to human performance, the team would assess appropriate metrics to measure performance and make the comparison. BAE developed a hardware platform with multiple sensors and an associated compute platform to perform automated fusion. This platform and associated testing is described below.

2.1   IMPLEMENTATION OF FUSION TEST PLATFORM

BAE developed an image acquisition suite for collection of coincident multi-modal imagery for fusion and metric evaluation. This platform consisted of RGB, SWIR broadband and LWIR broadband sensors mounted on a pan-tilt unit.

In addition to developing the sensor platform for a data collection to be used for fusion analysis, BAE integrated software components that were used to control the data collection and perform automated fusion and tracking. The sensor hardware and software components were fully integrated for the data collection in October 2008. Figure 1 shows a view of the sensor platform developed for testing, a block diagram of the software components, and the test site used for the data collection.
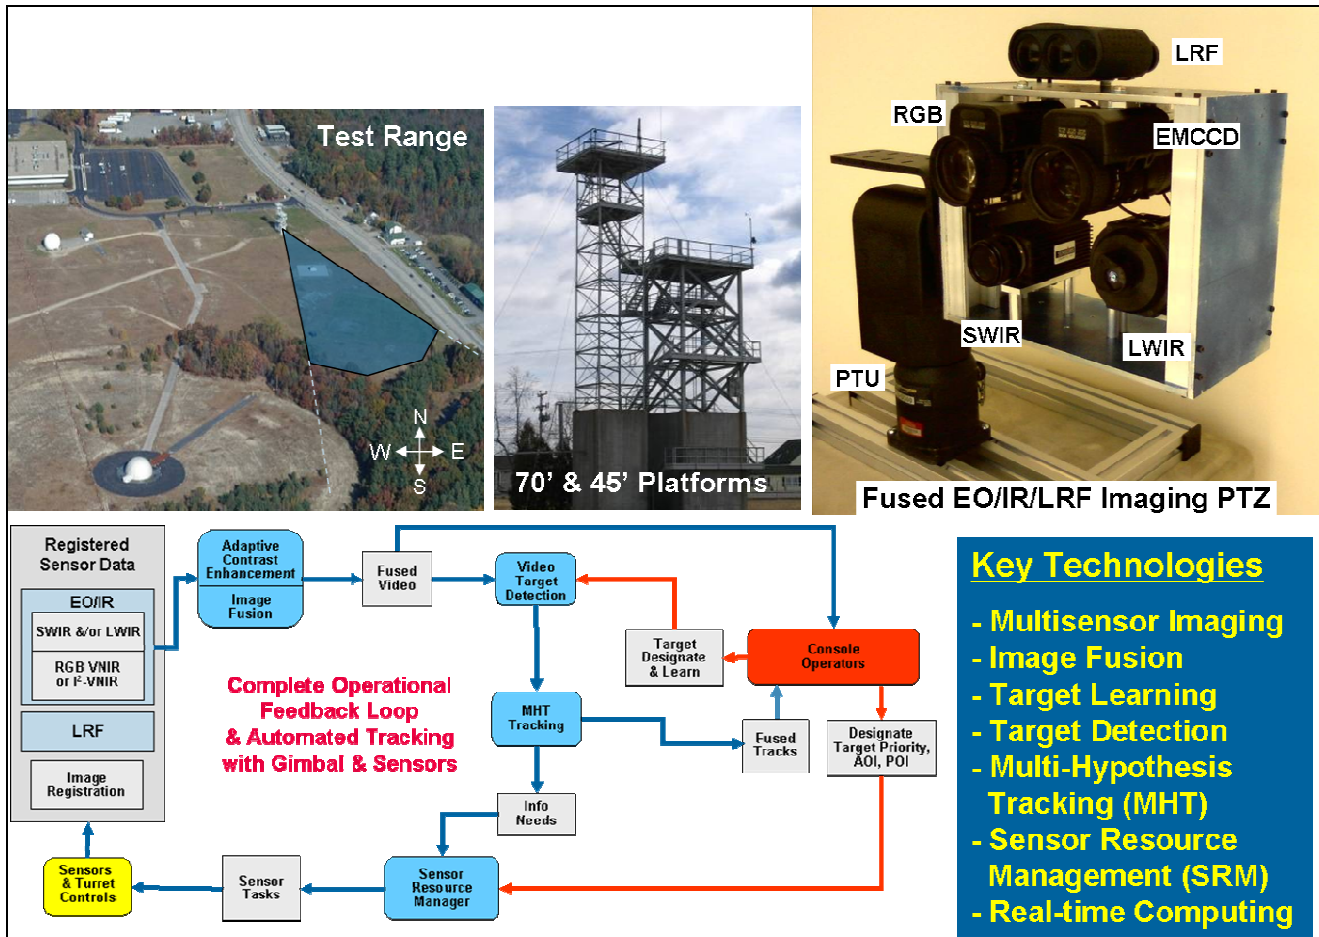
Figure 1. Hardware and software components integrated for fusion demonstration and data collection.

From the block diagram in Figure 1, it can be seen that imagery from the cameras underwent adaptive contrast enhancement and image fusion. The images were co-registered prior to these enhancement and fusion steps. Fused video was formed by applying neural color opponent visual processing, a concept demonstrated in Figure 2. In color opponent processing, local adaptive contrast and gain processing in individual image frames is followed by opponent color contrast processing, in which registered pixel values from multi-sensor image frames are contrast enhanced with resulting values mapped to RGB image layers. The contrast enhancement processing is achieved by center-surround shunting dynamics which enhance contrast peaks and valleys between multi-sensor image pixels. Center-surround processing is a neural processing step modeled after contrast enhancement processing in the human visual system. Image fusion processing consisted of the color opponent processing as well as a Fuzzy

ARTMAP classifier which adaptively trained models of object classes selected by the operator, enabling vehicles to be identified in subsequent video frames.

Figure 1 also shows that the system architecture included tracking processing in addition to the fusion algorithms. In the demonstration an operator nominated target vehicles via an interactive GUI, and the system detected nominated objects and tracked them in the video imagery through continuous motion. In addition, a sensor resource manager was used to drive control of the pan-tilt unit such that target vehicles were maintained in the camera fields of view. All of this was demonstrated with the sensor platform on top of the tower shown in Figure 1 and with two vehicles moving in crossing patterns on the test range shown in the figure.
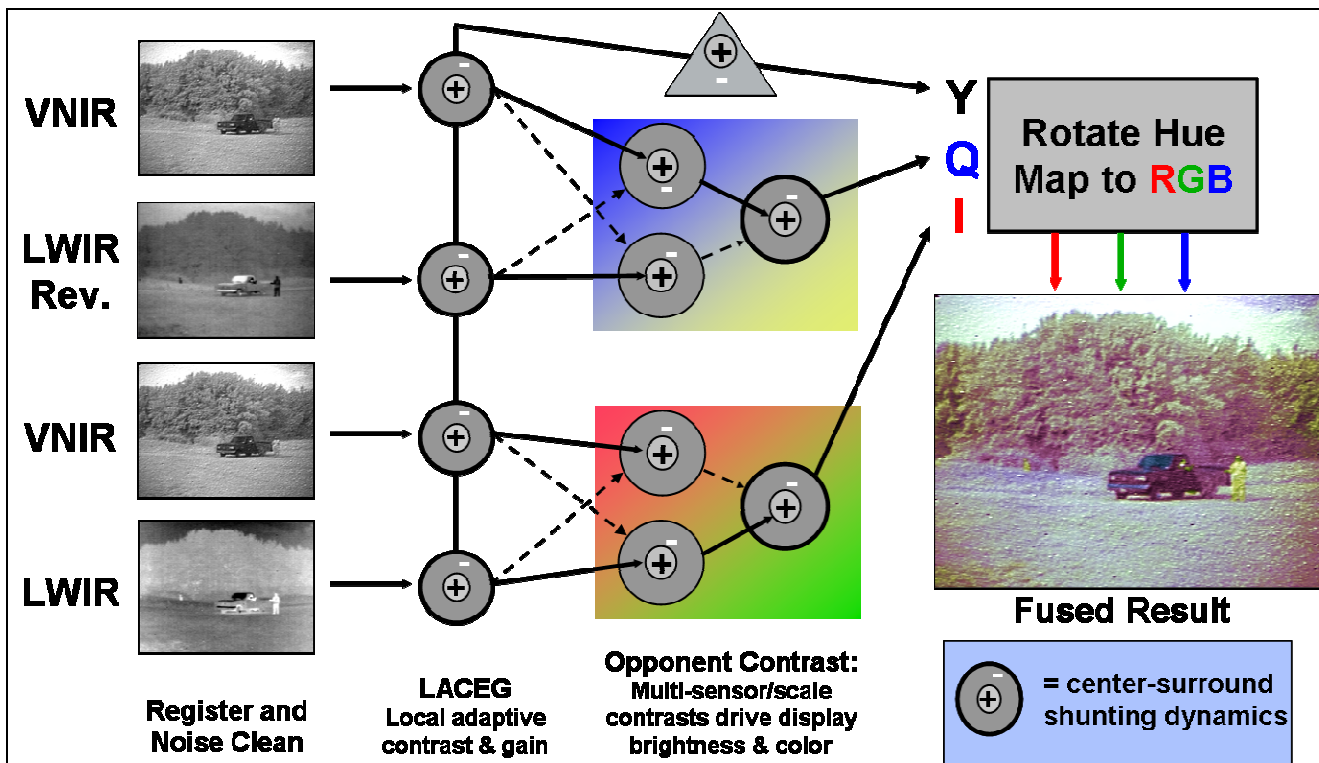

Figure 2. Dual-Sensor Image Fusion Architecture, Neural Opponent-Color Visual Processing Model

Figure 3 shows the real-time display from the demonstration system during the data collection. On the top left is the RGB camera imagery, on the top right the LWIR camera imagery, in the lower left the fused image formed from color opponent processing of the RGB and LWIR images, and in the lower right is the SRM/track display. The vehicles are distinct in all three image frames, and in particular the fused imagery shows a strong contrast between the vehicles and the rest of the background. In the track

display, track trails are shown for both vehicles. The expected field of view of the cameras is shown in the track display as well, as determined by the sensor resource manager which worked to optimize pointing of the pan-tilt unit to ensure that both vehicles were in the camera fields of view as much as possible.
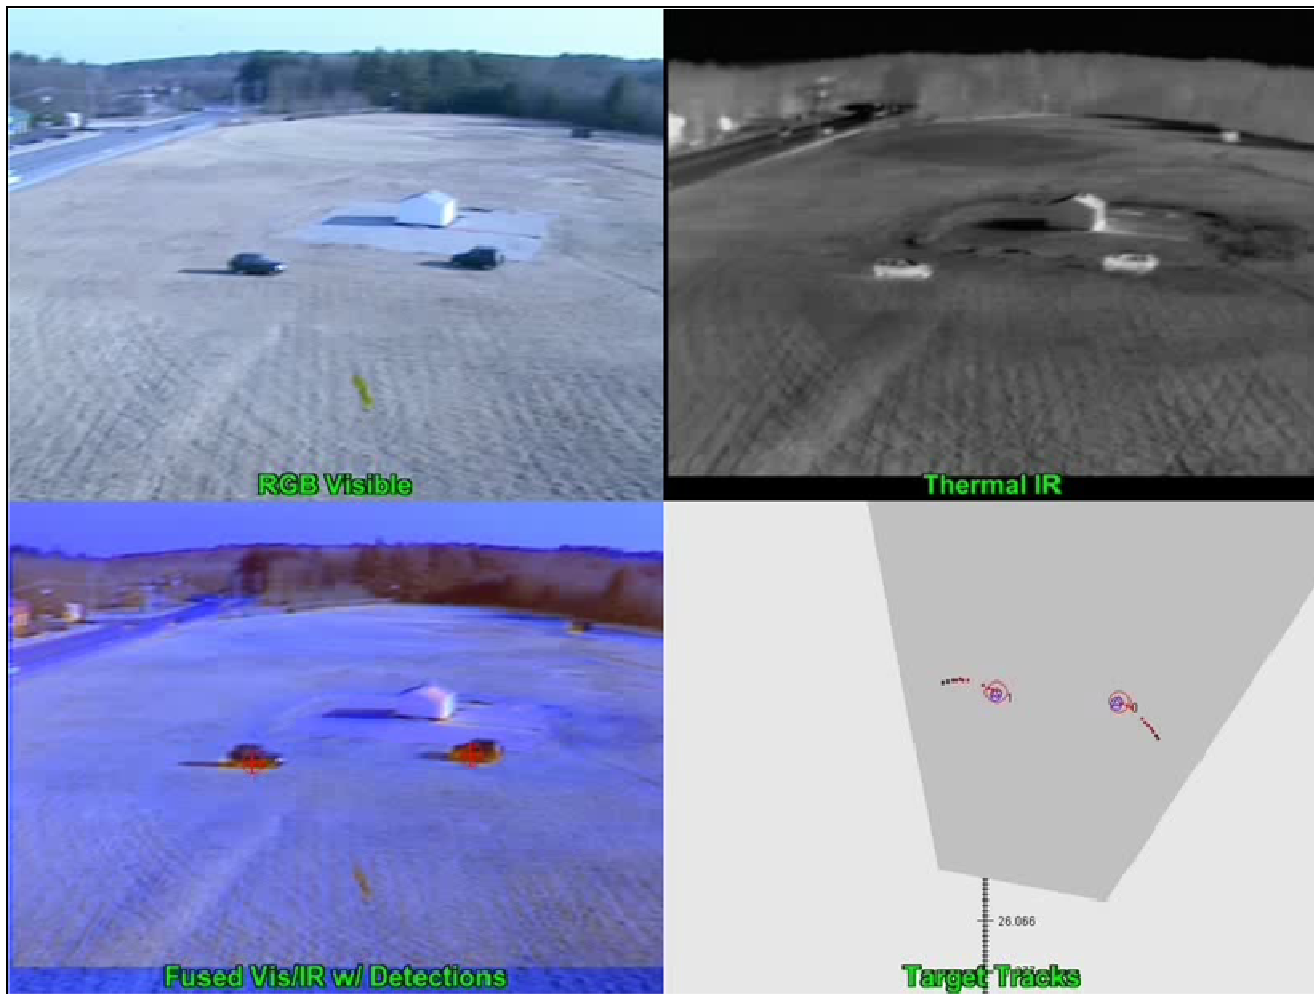

Figure 3. RGB/LWIR Fused Video Detection, Tracking & SRM Tower Demo

The fusion test platform successfully demonstrated real-time fusion of multi-sensor video, and associated tracking of fused detections. Initially the team intended to continue with this line of investigation to collect more data sets and to study metrics in evaluating fusion processing of this data as compared to human analysts. The research team decided to focus instead on using the same fusion algorithms within a software testbed to perform fusion metric testing and evaluation. Those studies are described in the next section.

3 FUSION METRICS SOFTWARE TESTBED

The final year of the Fusion Metrics program focused on development of a software test platform in which fusion algorithms could be studied in terms of quantitative performance metrics. While fusion algorithms inherently must be tested via some means of software testing, the focus in this research was on creating a modular software environment in which feature vectors from any new data set could be placed in a database and fusion testing could be performed using any modular fusion component that complies with the testbed interface. Monte Carlo testing of the fusion algorithms would then produce statistical performance results that could be analyzed using a variety of metrics. In this research, the team created the software testbed and performed metrics testing using an AFRL dataset.

3.1 FUSION SOFTWARE TESTBED DESIGN

The focus of designing the software fusion analysis testbed was on creating a modular design that enabled rapid fusion prototyping and performance analysis. Figure 4 shows a block diagram of the software components in the testbed. Like the previous work in the Fusion Metrics program which focused on image fusion, this testbed was tested with video data from AFRL. Similar to the dataset collected by BAE in tower testing described in Section 2, the AFRL data included EO and IR video. The feature layer processing was similar to that used in the BAE sensor fusion demonstration system described in Section 2, where the use of color opponent processing and normalized RGB and IR layers was the main method of producing feature layers. The video included vehicles, so once again the effort focused on extracting features associated with vehicle pixels in the imagery. Each pixel in a given image had an associated vector stack consisting of the image feature layers formed from the original RGB and IR imagery, and for all pixels associated with a vehicle these feature vectors were stored in a database. All feature processing in the software testbed was performed with Matlab code and the resulting features were logged in a database.
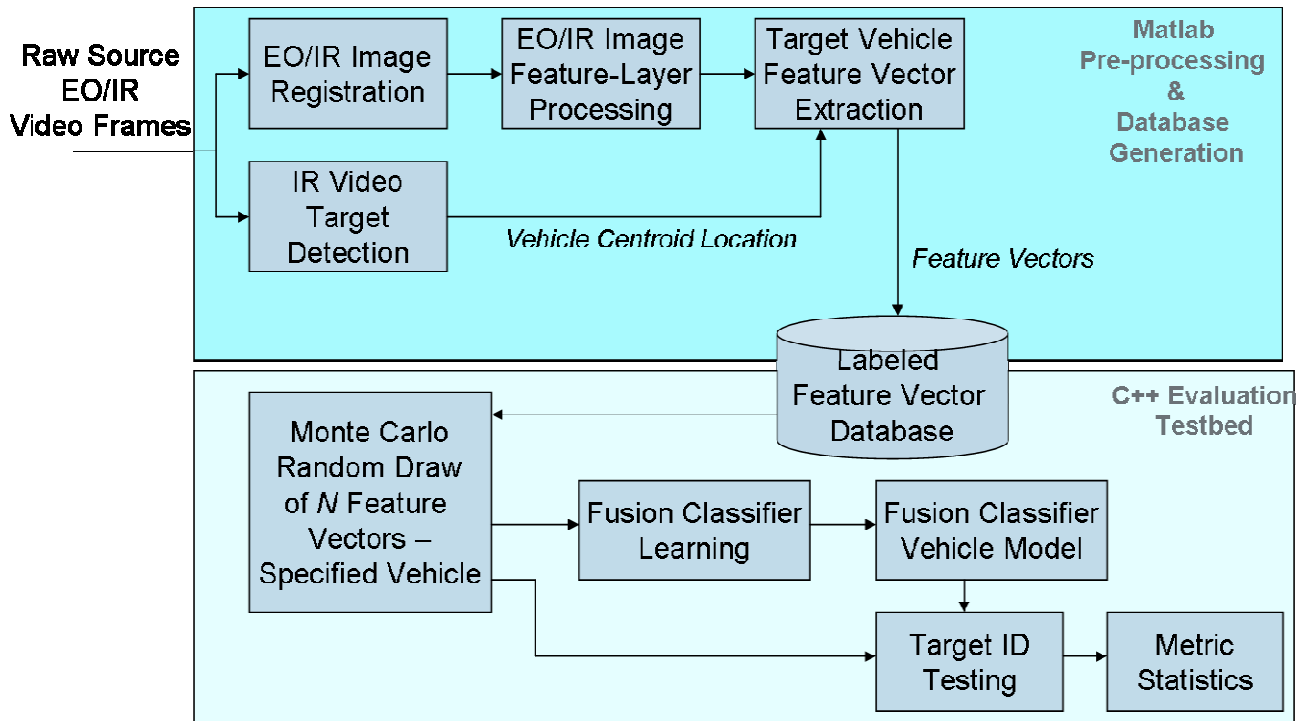
Figure 4. Software Testbed Design for Evaluation of Fusion Metrics

The Matlab code for feature extraction and database storage represented the front-end of the software fusion testbed. The back-end was a C++ based Monte Carlo test engine that randomly extracted feature vectors associated with specific vehicles and tested the fusion classifier component of the system for identification performance. Like the system described in Section 2, this fusion processing utilized a Fuzzy ARTMAP algorithm. Figure 5 shows the fusion processing chain used in this system, in which opponent-color processing produced multi-sensor image features that were learned in the Fuzzy ARTMAP neural network. The result of the training step is a representation in feature space of the target and possibly counter-example vehicle. While this was the fusion methodology demonstrated in this software testbed implementation, the system is modular such that another fusion algorithm could be dropped in for similar performance analysis.

The Monte Carlo tests were performed in two stages – a learning stage and a testing stage. In the learning stage, feature vectors from a specific target vehicle and possibly a confuser vehicle were used to create a fusion model of the target. In the testing stage, feature vectors from the true target and from confusers were tested against the model and the system produced a target association for each test vector. Metrics were stored on these outcomes that could then be used for fusion performance and metric analysis.
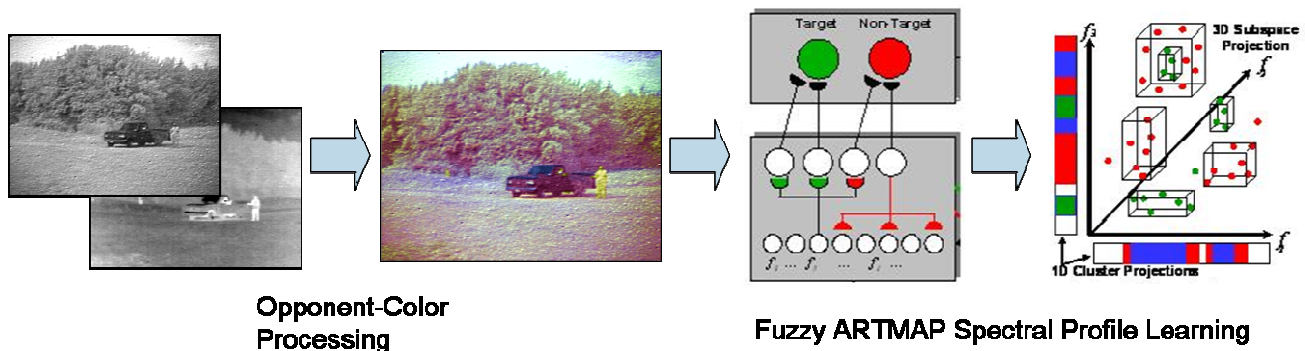
Opponent-Color
Processing

Fuzzy ARTMAP Spectral Profile Learning

Figure 5. Video fusion achieved via opponent-color processing and Fuzzy ARTMAP neural network.

3.2   FUSION METRICS PARAMETER ANALYSIS

In these studies, a metric that showed particular utility was the Probability of Correct Association (P_{CA}), which is the rate at which true target vehicles and confuser vehicles were correctly identified during the testing stage. The metric could be studied against various free parameters to enable the most thorough analysis of fusion results. For example, Monte Carlo tests were run for different sets of feature vectors extracted from the imagery. Feature layers were formed from a variety of color opponent combinations and RGB and IR layer normalization methods. Then, different groupings of the feature layers were formed such that each different combination represented a unique set of image features. Fusion results were not just dependent on the fusion algorithm but on the specific features extracted. The fusion metrics testbed enabled researchers to rapidly assess the utility of a variety of feature sets in the context of the fusion algorithm under test.

Figure 6 shows an example outcome from the fusion metrics testbed using toy data. This example demonstrates the type of analysis that was possible, where three different feature sets were compared for P_{CA} in fusion testing. The P_{CA} outcomes are compared based on a variable number of testing frames but for a fixed number of training frames. This type of analysis would enable an analyst to rapidly determine the feature set that offered the most utility in the fusion system for a given data set, and to determine the best operating point for both training and testing set size.
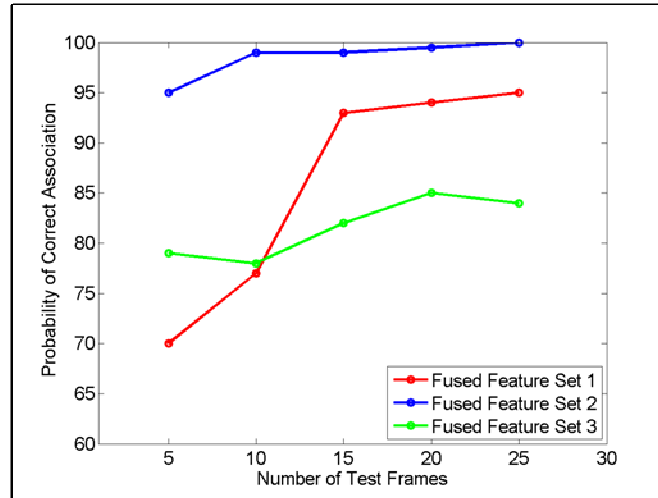
Figure 6. Example analysis outcomes from Fusion Metrics testbed.

  Figure 6 represents one kind of analysis outcome, but many others could be performed. There are a variety of free parameters that could be examined – comparison of fusion outcomes for different target types, for comparison of fusion algorithm parameters to find optimum fusion parameters, to compare the durability of fusion processing across data sets which may have been collected at different times, or to choose different fusion performance metrics than P_{CA} and to determine viability of the metric. In the analysis of the AFRL dataset, the key focus was on determining optimum feature layers and understanding how fusion performance decayed with respect to time differences between the collection of feature sets for learning versus testing.

4 CONCLUSIONS

The Fusion Metrics program ultimately demonstrated two unique research capabilities. The first, described in Section 2 and serving as the focus of the first two years of the program, demonstrated a capability to fuse multi-sensor imagery in real-time with associated tracking algorithms. But more importantly, this platform demonstrated a method to rapidly collect image data that could be used for fusion studies. Often times the lack of relevant datasets can hinder fusion studies, and in this platform BAE developed a data collection platform that could be used in a variety of environments to collect useful datasets for fusion algorithm analysis and metric evaluation.

The Fusion Metrics software testbed was the focus of the final year of the program and demonstrated a modular test platform design that enabled rapid prototyping of a fusion algorithm on video datasets. Section 3 showed an example data analysis outcome, a P_{CA} plot that compared fusion performance for three distinct feature sets. This was an example of one type of metric analysis, but the software platform enables fusion metric analysis against the many free parameters that affect fusion outcomes. Having a test platform that enables such rapid prototyping may prove critical for any research focused on optimization of a fusion system, or perhaps an initial understanding as to whether a given fusion algorithm is likely to perform at all. The modular design includes a database that could be filled with any useful feature set. This capability would enable a researcher to assess whether a fusion algorithm is robust against numerous types of feature data, or whether is can only be expected to perform for distinct classes of data.

Ultimately, the analysis of fusion metrics must continue in the future. There is likely no single metric that can optimally describe fusion performance for all classes of algorithms or data. Therefore, it is important to evaluate fusion with a meaningful metric or metrics and to study the performance against many relevant parameters that affect fusion outcomes.

5 REFERENCES

1. **Active Tracking of Surface Targets in Fused Video**, Allen Waxman, David Fay, Paul Ilardi, Pablo Arambel and Jeffrey Silver, *Proceeds of 10th Int'l Conf on Information Fusion,* Fusion 2007, Quebec

2. **Sensor Fused Night Vision: Assessing Image Quality in the Lab and in the Field,** Allen Waxman, David Fay, Paul Ilardi, Dick Savoye (consultant), Robert Biehl, and Dan Grau, *Proceeds of 9th Int'l Conf on Information Fusion,* Fusion 2006, Florence

3. **Realtime Image Fusion and Target Learning & Detection on a Laptop Attached Processor,** David Fay, Paul Ilardi, Nathan Sheldon, Daniel Grau, Robert Biehl and Allen Waxman, *Proceeds of Enhanced & Synthetic Vision* SPIE Vol. 5802, 2005

4. **Fusion of Visible and Thermal Imagery Improves Situational Awareness**, A. Toet, J.K. IJspeert, A. Waxman, and M. Aguilar; *Proceeds. of Enhanced & Synthetic Vision 1997,* SPIE-3088, 1997; also *Displays,* 18, 85-95, 1997