United States Air Force
Scientific Advisory Board

Report on

Automatic Target Recognition

Executive Summary and Annotated Brief (PR)

SAB-TR-05-02-PR
July 2005

DISTRIBUTION AUTHORIZED
In accordance with AFI 61-204 and DODD 5230.24, distribution statement A, this document is approved for public release; distribution is unlimited.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
01 JUL 2005	**Final**	**01 JAN 2004 - 01 JUL 2005**

4. TITLE AND SUBTITLE
Automatic Target Recognition, Executive Summary and Annotated Brief (PR)

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Dr. Llewellyn Dougherty

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
HQ USAF/SB 1180 AF PENTAGON RM 5D982 WASHINGTON, DC 20330-1180

8. PERFORMING ORGANIZATION REPORT NUMBER
SAB-TR-05-02-PR

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
SAF/OS, AF/CC AIR FORCE PENTAGON WASHINGTON, DC 20330-1670

10. SPONSOR/MONITOR'S ACRONYM(S)

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

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT
The complexity of warfare and the requirement to shorten the kill chain against difficult targets has increased the need for Automatic Target Recognition (ATR). No single sensor currently provides a robust enough capability of detecting all classes of targets through the camouflage and deceptive elements of natural and/or man-made clutter. ATR is a sub-element of a larger class of problems known collectively as Target Identification. This study concentrated on image-based techniques as an ATR subdiscipline of target ID. The study panel concluded that no silver bullet exists that will dramatically improve ATR performance in the near term. However, more work and understanding in very specific areas could improve the way the Air Force uses ATR today. For example, CONOPS and technology can work synergistically (e.g. change detection) to improve mission capabilities; ATR can be an aid to, but not a substitute for, human judgment by both pilots and image analysts; and ATR is essential to better match image exploitation capabilities with collection capacity.

15. SUBJECT TERMS
Automatic Target Recognition, ATR, Sensor, SAR, Synthetic Aperature Radar, Change Detection, Multi-Sensor, Wide Area Search, ISR, Persistence, Attack

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**SAR**	**62**	

United States Air Force

Scientific Advisory Board

Report on

Automatic Target Recognition

Executive Summary and Annotated Brief (PR)

(This Page Intentionally Left Blank)

Foreword

The complexity of warfare and the requirement to shorten the kill chain against obscured, difficult, and lethal targets has increased the need for Automatic Target Recognition (ATR). No single sensor currently provides a robust capability to detect all classes of targets through camouflage and the deceptive elements of natural and/or man-made clutter. Combat target identification is further degraded by other operational complications such as urban context, electronic environment, mobile targets, weather, and the presence of bomb damage debris.

ATR technology has been around for a long time. Optical and electronic image correlators date back at least to the early 1960's as tools for registering two images or for locating one image within a larger image. However, previously fielded systems fall short as robust solutions for accurately detecting, identifying, and tracking both stationary and mobile targets. The system applications were narrowly focused, had limited objectives, and have all been through repeated cycles of development, improvement and technology insertion.

This study addresses the challenge of defining a path to accelerate ATR to the warfighter. The study was conducted in response to a request by the Secretary of the Air Force and the Air Force Chief of Staff.

The ATR Quick Look study team conducted a series of briefings and interviews with key developers and users of ATR systems. These individuals included system developers, operational commands, and various research agencies from the Air Force, DOD and industry. The assistance of these organizations was essential to the completion of the study and guided the SAB team toward their findings and recommendations. The study team greatly appreciates the cooperation of these organizations, and acknowledges their valuable contributions.

The undersigned also wish to acknowledge the outstanding effort put forth by the other members of the ATR study team and the Air Force Scientific Advisory Board Secretariat – whatever value is found in this work is attributable to their tireless efforts.

Dr. Llewellyn Dougherty
ATR Study Chair

(This Page Intentionally Left Blank)

Executive Summary

Introduction

Automatic Target Recognition (ATR) is a sub-element of a larger class of problems known collectively as "Target Identification". There are a number of sub-disciplines within Target Identification that include elements of ATR. Some are based on one-dimensional signal processing (e.g., high range-resolution radar), some are based on multi-dimensional and perhaps multi-sensor data fusion, and some are based on behavioral or dynamic response characteristics. Furthermore, some are based on exploitation of two-dimensional imagery such as electro-optical (EO), infrared (IR), and synthetic aperture radar (SAR). This study concentrated on image-based techniques.

Image-based ATR can be applied to many phases of the effects chain. ATR is relevant to Find, Fix, and Track -- the parts of the chain that we refer to as intelligence, surveillance, and reconnaissance (ISR). ATR is also pertinent to Engage, which is most relevant to manned attack systems and weapons. For this reason, the Study focused predominantly on the ISR aspects of ATR when ATR is applied as aids for image analysts (IAs), searching for mobile "targets", through-the-weather, and using radar image products.

In short, the study did not discover solutions that would radically advance ATR to a truly *automated* target recognizer. However, there are significant things ATR can do to assist IAs and improve mission capabilities.

Study Purpose

At the request of the Secretary of the Air Force and the Air Force Chief of Staff, the Air Force Scientific Advisory Board (SAB) undertook a quick look study of Automatic Target Recognition technology for fiscal year 2005. A small team was assembled from current SAB members, augmented by a small number of outside experts, including participants from the Air Force Research Laboratory Sensors Directorate (AFRL/SN). The Study solicited and received presentations from a wide variety of government, industry, and Federally Funded Research and Development Center (FFRDC) experts. Specifically, the team was tasked to:

1. Assess
    - Advances in modeling & simulation capabilities, signature libraries, and exploitation tools;
    - Progress and effectiveness in state-of-the-art science and technology (S&T) in ATR algorithms under different environments;
    - Challenges of ATR and conditions which may limit the effectiveness of ATR;
    - How far the AF can move from *Assisted* Target Recognition to fully *Automated* Target Recognition; and
2. Define a path to accelerate ATR to aid the warfighter.

State of ATR

The study panel concluded that no "silver bullet" exists that will dramatically improve ATR performance in the near term. However, more work and understanding in very specific areas could improve the way the Air Force uses ATR today. For example, Concept of Operations (CONOPS) and technology can work synergistically (e.g. change detection) to improve mission capabilities; ATR can be an aid to, but not a substitute for, human judgment by both pilots and IAs; and ATR is essential to better match image exploitation capabilities with collection capacity.

ATR technology is worthy of continued investment. The study panel believes investment in the following areas would realize firm gains in ATR:

- Operational and advanced sensor image libraries
- Algorithms, modeling, and simulation-based evaluation
- Independent algorithm and system testing
- Cognitive human-machine interface for users
- Improved operator/technologist collaboration

Recommendations for Improving ATR Application and Development

The study made five recommendations to the Air Force for improving how ATR is currently applied and developed:

1. **ATR as a System.** The Air Force should assure future ATR maturation efforts are treated in a system context, to include sensors, users, and CONOPS, as well as ATR.
2. **Network Operators and Technologists.** A need exists to formalize and implement a mechanism through which the operators and technologists are networked together (on both SIPRNET and JWICS). The network will enable shared data, cooperative development, and experimental validation of ATR capabilities. Include, at a minimum, a DCGS/DGS site, DMOC, and AFRL's Sensors and Information Directorates.
3. **Certify ATR Systems.** An independent activity and process (similar to the role of the Central Inertial Guidance Test Facility (CIGTF) for inertial navigation systems) should be established to characterize and certify ATR systems -- both "automated" and "assisted" -- prior to acquisition.
4. **Target Identification Study.** Ground target identification, of which ATR is a small subset, proves to be a difficult albeit important problem. In order to better assess ground target identification capabilities, progress, and technology needs, a comprehensive study should be initiated in this area.

5. **Continue Investment in ATR.** The Air Force should support ATR-related technology development with continued investment and management attention. Of particular note, investment in the following areas would be highly beneficial:
   - Operational and advanced sensor image libraries
   - Algorithms, modeling, and simulation-based evaluation
   - Independent algorithm and system testing
   - Cognitive human-machine interface for users

(This Page Intentionally Left Blank)

Table of Contents

(This Page Intentionally Left Blank)

Section 1: Introduction

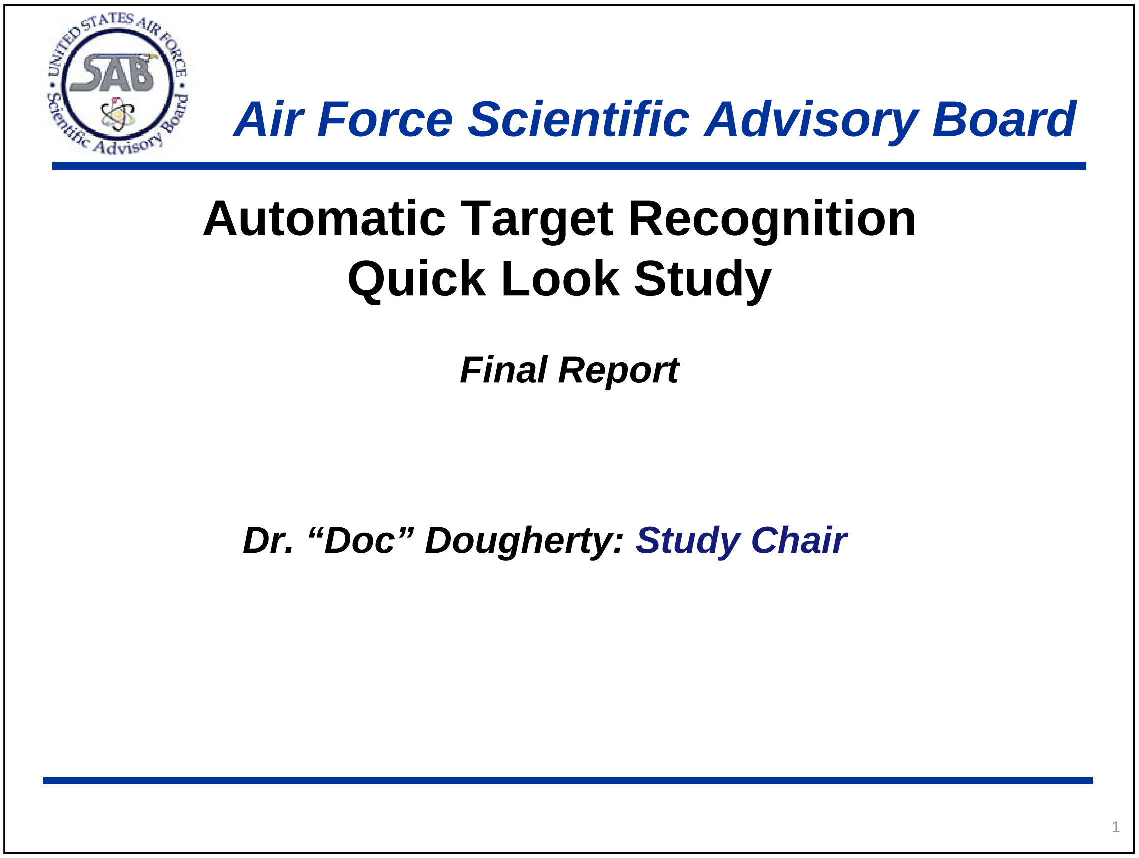

At the request of the Secretary of the Air Force and the Air Force Chief of Staff, the Air Force Scientific Board (SAB) undertook a quick look study of Automatic Target Recognition (ATR) technology in the Spring of 2004.

Panel Members and Staff

SAB Members and Consultants

Dr "Doc" Dougherty	Raytheon
Dr Bob Selden	Los Alamos (Retired)
Dr Lou Metzger	MITRE
Mr Scott Fouse	ISX
Dr John Betz	MITRE
Dr Jurgen Gobien	SAIC
Dr Pete Worch	Consultant

General Officer Participant

Brig Gen Paul Capasso	Director, C4 USTRANSCOM

Government Advisors

Col Dan DeForest	Mil Asst to AF Ch Scientist
Ms Lori Westerkamp	AFRL/SN
Mr Ed Zelnio	AFRL/SN
Dr Mike Eismann	AFRL/SN

SAB Staff

Maj Pia Lomax	Executive Officer
Maj Kyle Gresham	Military Assistant
Capt Perry Villanueva	Technical Writer
Mr Jay Carlson	Technical Assistant

3

A small team was assembled from current SAB members, augmented by a small number of outside experts, including participants from the Air Force Research Laboratory Sensors Directorate. Brigadier General Paul Capasso, US Transportation Command (USTRANSCOM) Director of Command, Control, Communication, and Computers (C4), served as the General Officer Participant. The SAB provided invaluable support through a combination of organic staff, contract support, and personnel loaned from other Air Force organizations.

Presentations

- DARPA
- SAIC
- AFRL/IF
- AFRL/MN
- AF/XOI
- BAE
- The Aerospace Corporation
- NGA
- 480 IG
- Raytheon
- Northrop Grumman
- Johns Hopkins University
- Scitor Corp (NRO)
- MITRE (DCGS Program Office)

4

The Study solicited and received presentations from a wide variety of government, industry, and FFRDC experts, including those listed above. A total of nine days of fact-finding was conducted over three months between January and March 2005.

Terms of Reference

- **Assess**
  - **Advances in modeling & simulation capabilities, signature libraries, and exploitation tools.**
  - **Progress and effectiveness in state-of-the-art S&T in Automatic Target Recognition algorithms under different environments.**
  - **Challenges of ATR and conditions which may limit the effectiveness of ATR.**
  - **How far the AF can move from Assisted Target Recognition to ATR [fully Automated Target Recognition]**
- **Define path to accelerate ATR to aid the warfighter.**

5

The complete Study Terms of References are included in Appendix A and summarized here.

Section 2: Background

Sub-disciplines of Target Identification

- **Target Identification involves many technology pieces:**
  - **Image-based automatic target recognition**
  - **Multi-sensor fusion**
  - **High range-resolution radar**
  - **Radar moving target tracking**
  - **Combat ID**
  - **Intelligence data correlation**
  - **Motion exploitation**
  - **Signatures and Modeling**

This Quick-Look Study concentrated on image-based ATR

6

ATR is a sub-element of a larger class of problems known collectively as "Target Identification". There are a number of sub-disciplines within Target Identification that include elements of ATR. Some are based on one-dimensional signal processing (e.g., high range-resolution radar), some are based on multi-dimensional and perhaps multi-sensor data fusion, some are based on behavioral or dynamic response characteristics, and some are based on exploitation of two-dimensional imagery (EO, IR, and SAR). This quick look study concentrated on the image-based techniques.

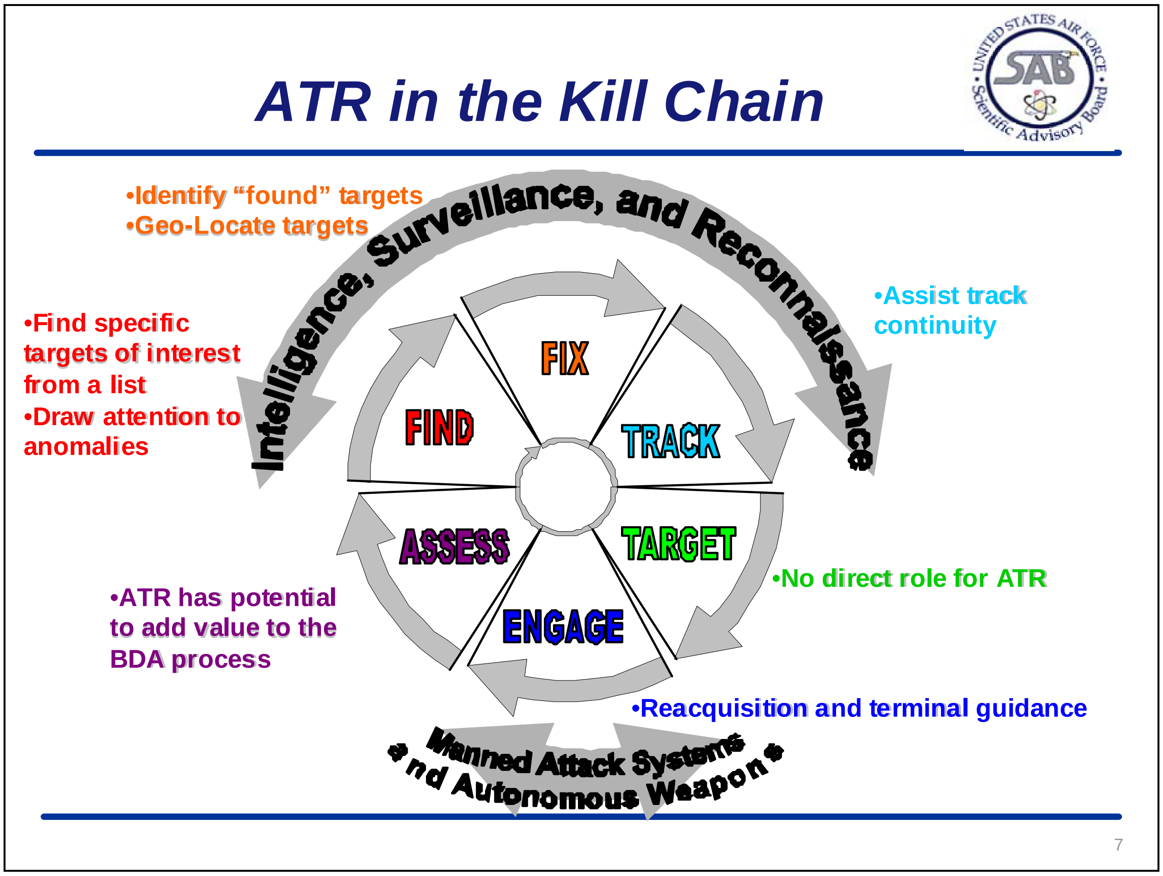

Image-based ATR can be applied to many phases of the effects chain. ATR is relevant to Find, Fix, and Track (the parts of the chain that we refer to as ISR) as well as to Engage (most relevant to manned attack systems and weapons). This Study focused predominantly on the ISR aspects of ATR when used as:

- Aids for IAs
- Searching for mobile "targets"
- Through-the-weather
- Using radar image products

It should be noted that in this context Targeting is strictly a human decision process that designates an object (perhaps selected by an ATR system) as a target and selects the particular effect to be rendered.

In many respects, the application of ATR to Engage is a repeat of its application to Find, Fix, and Track, but only after the Targeting decision has been made. Thus when applied to Engage, the search area has been refined to a point where false alarms are unlikely and the ATR can be set to achieve a high detection probability.

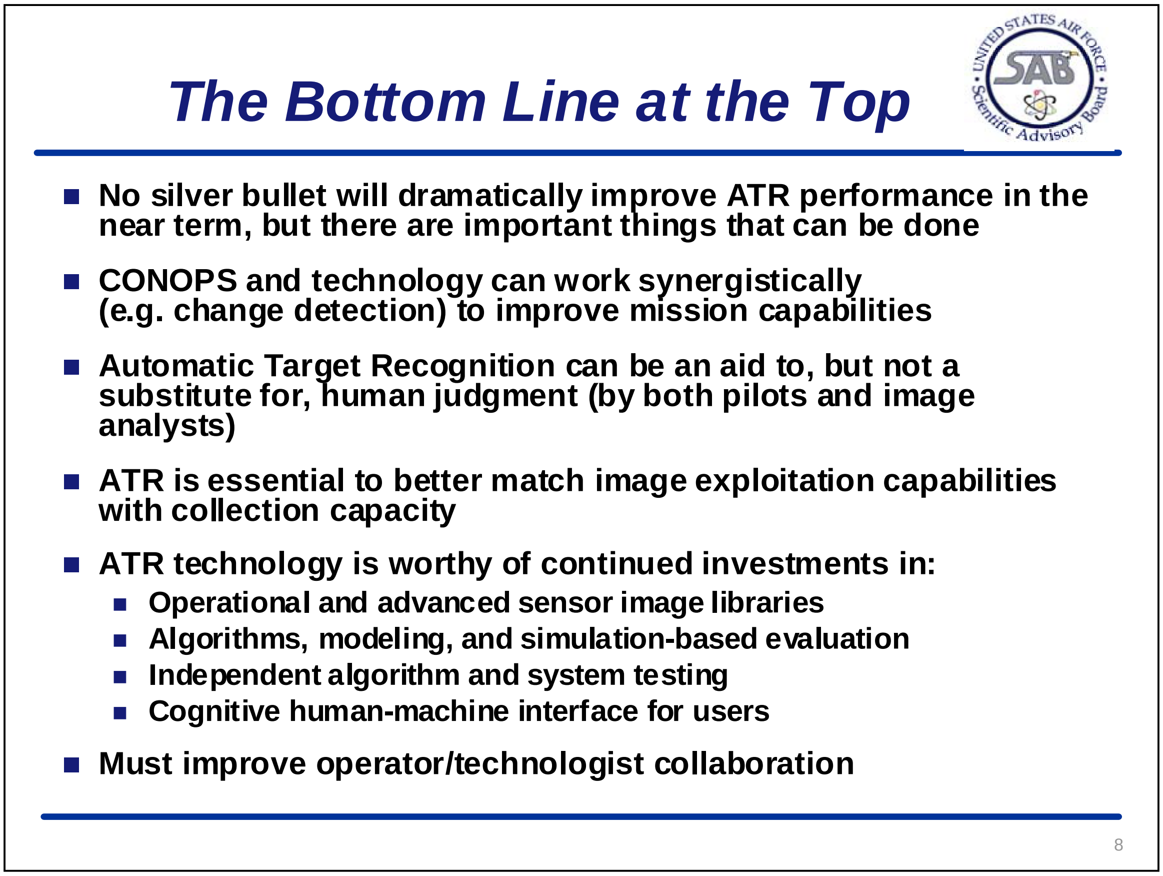

For those who are in a hurry, we present here the short version of the Study conclusions. The remainder of this document provides both tutorial material and more detail, including

- Some ATR history
- What ATR does
- Characterization of ATR performance
- Details of
    - Six findings, and
    - Five recommendations.

Historical Context

- Long (30+ years) history of investment in ATR
- Many attempts to transition ATR into systems
- The few successful transitions
  - Tomahawk & ALCM (SMAC/TERCOM used for INS update only)
  - Longbow Apache (Ka-band real beam radar)
  - Air-to-Air ID (multiple techniques)
- All share some common attributes:
  - Narrowly focused
  - Limited objectives
  - Iterative development

9

ATR has been around for a long time. Optical and electronic image correlators date back at least to the early 1960's as tools for registering two images or for locating one image within a larger image.

TERCOM (TERrain COntour Matching) is an ATR technique that uses the uniqueness of terrain altitude profiles to update cruise missile navigation systems based on radar altimeter measurements. It was developed in the mid 1960's and was first fielded in the Air Launched Cruise Missile (ALCM) in the 1980's.

The optomechanically implemented Scene Matching Area Correlator (SMAC) and its digital grandchild (DSMAC) originated at the Naval Avionics Facility Indianapolis in the late 1960's and was fielded in Tomahawk beginning in the 1980's.

The ATR sub-system fielded as a part of the Apache-Longbow radar was in development for over 10 years prior to its production delivery as part of the operational weapon system.

These fielded ATR systems share some common attributes:

- Each application was narrowly focused
- Each program had limited objectives, and

- All have been through repeated cycles of development, improvement and technology insertion

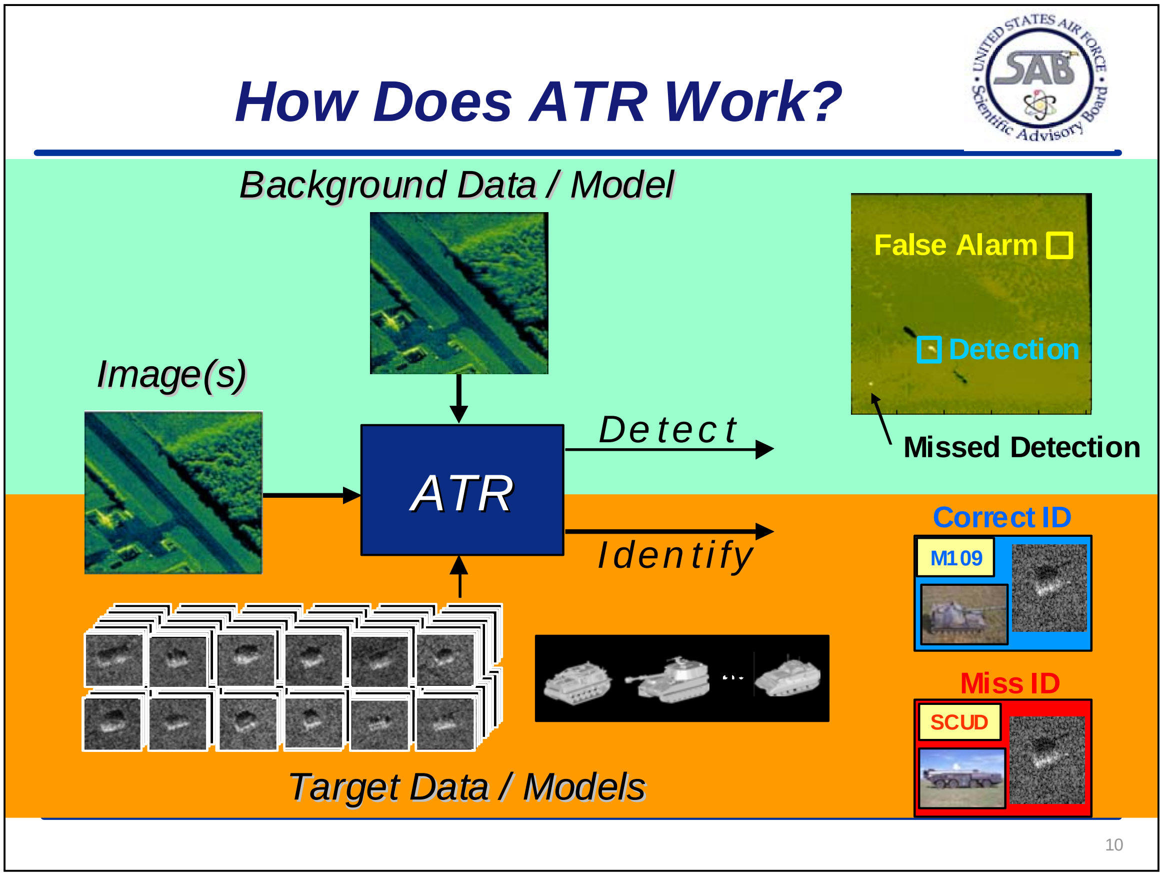

Automatic Target Recognizers (ATRs) process sensor data to detect and identify targets. ATRs make detection and identification (ID) decisions by comparing the incoming sensor data to a pre-stored database. The database can store both background information and target signature information.

Change detection, for example, compares the incoming sensor data to pre-stored background data to detect changes (e.g., targets moving in or out of an area). For change detection, the background database requires at least one previous image of the area.

ID, on the other hand, requires a large database to enumerate the various signatures of a target as a function of pose, configuration, background, etc. The performance of an ATR is directly related to the completeness of the database.

The detection performance is characterized by measuring the correct detections, the missed detections, and the false alarms. These measures are fundamentally dependent on each other as will be described in the next chart.

The ID performance is characterized by measuring the ratio of correct IDs to missed IDs. A correct ID labels a mission target (a target type contained in the database) as the correct target. A missed ID either 1) incorrectly labels a mission target as the wrong target or 2) incorrectly labels a confuser object as a mission target.

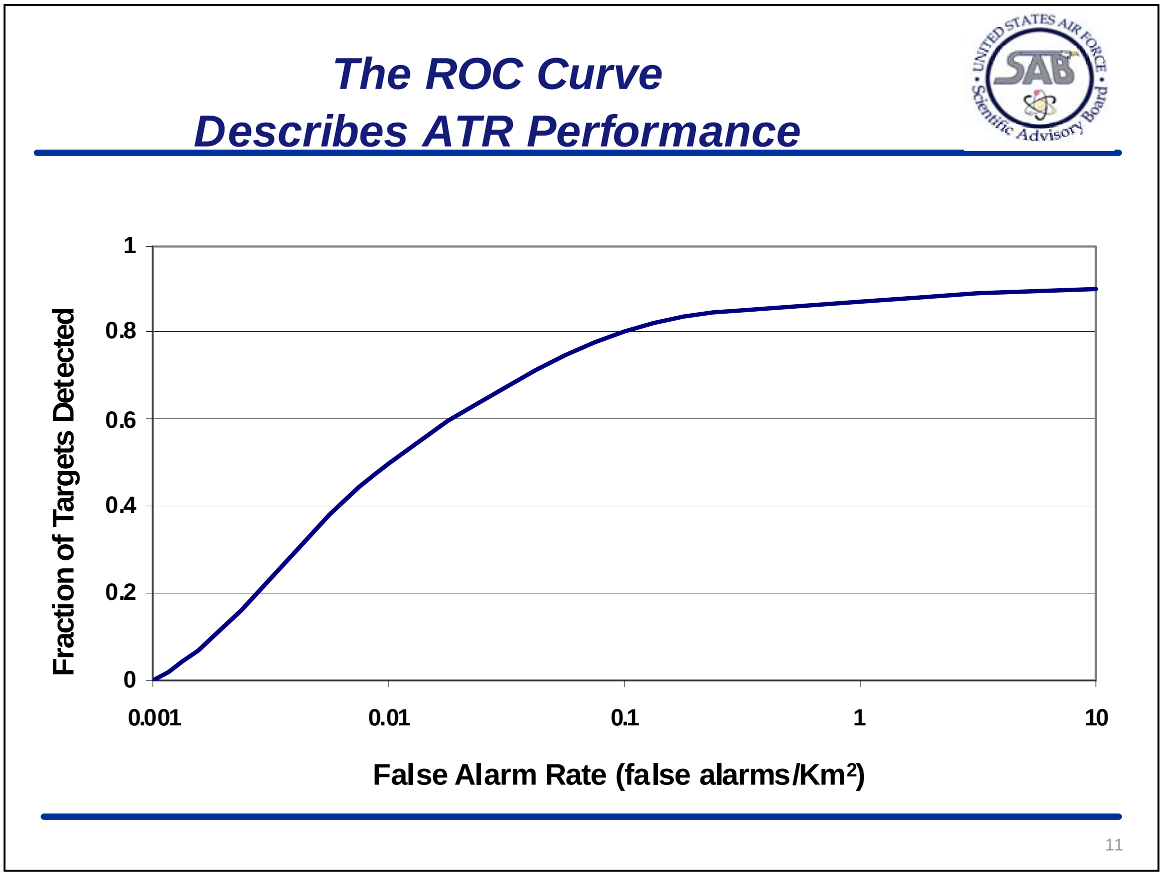

The Receiver Operating Characteristic (ROC) curve describes the range of performance options that are available from a particular ATR system under a particular set of circumstances. The ROC reflects both the statistics of the "target" and the statistics of the background. The developer (and in well designed systems, the user) selects the operating point.

For an example of how a ROC curve is derived see Appendix E.

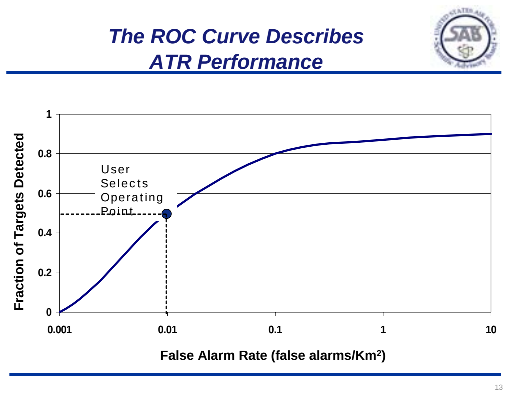

The selection of the operating point determines BOTH the detection probability AND the false alarm rate.

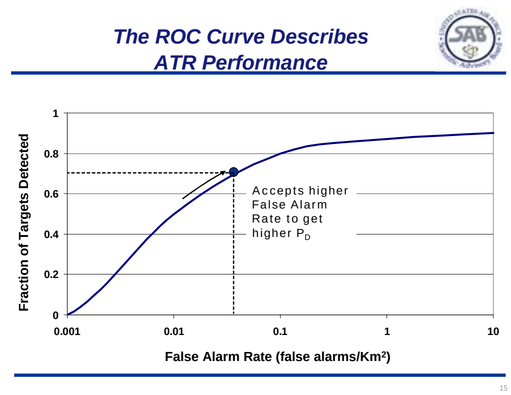

It is also possible to select an operating point that yields a higher detection probability, with a correspondingly higher false alarm rate.

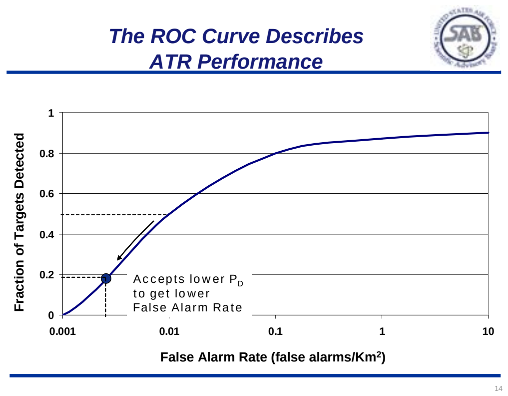

It is possible to select (elect) a lower detection probability in order to achieve a lower false alarm rate (as in cases where post-detection processing capacity limits total system throughput).

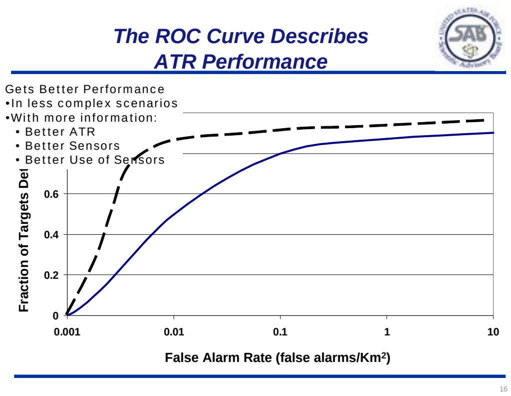

Under less stressing circumstances, a particular ATR will perform “better” in the sense that the ROC will lie above the more stressing case (the detection probability will be greater for the same false alarm rate).

Performance improvements might result from better ATR algorithms, better sensors, or better use of sensors.

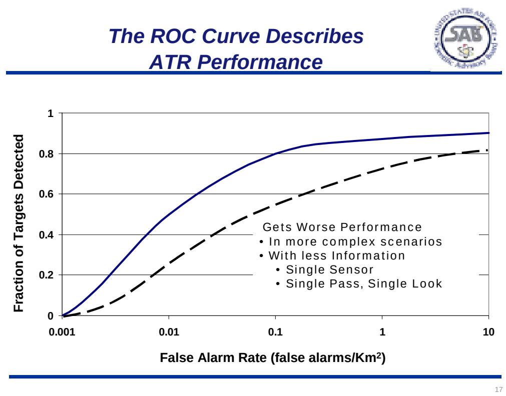

Performance may deteriorate in more complex scenarios or when sensors are limited by weather, geometry, or other factors. Under these circumstances the ROC curve may fall below more nominal performance values.

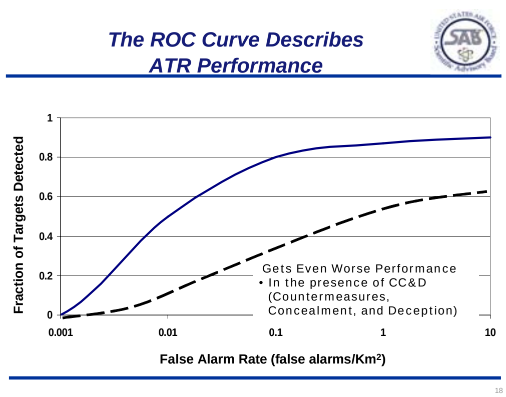

In the presence of countermeasures, concealment, and deception (CC&D), ATR performance can be expected to further deteriorate as the false alarm rate associated with a particular detection probability rises.

(This Page Intentionally Left Blank)

Section 3: Findings

Findings: ISR #1

- **Current operational sensors are quite capable of overloading image analysts**
- **Existing ATR technology and CONOPS are inadequate to address this problem**
- **A combination of ATR technology, new CONOPS, and new sensors may be required to balance future collection capabilities with anticipated exploitation capacity**

19

To deal with this bottleneck, we believe that tasking doctrine (constrained by exploitation capacity) limits collection by using less than full sensor capability.

Existing ATR technology and CONOPS are inadequate to address this problem.

Current False Alarm Rates (FAR) for Single-Look, Single-Pass, 1-meter SAR ATR are as much as two orders of magnitude too high to support wide area screening of images produced by one Global Hawk at maximum area coverage rate. Attempting to use existing technology to screen large imaged areas will produce more detections and false alarms (combined) than an analyst can service in the time available.

A combination of ATR technology, new CONOPS, and new sensors may be required to balance future collection capabilities with anticipated exploitation capacity.

As Global Hawk CONOPS are being developed, tasking has been constrained to not exceed exploitation capacity with the result that current collection rate represents a small fraction of sensor capability:

- “Occupancy check” is still widely employed to efficiently use available analysts
- TPED (Task, Process, Exploit, and Disseminate) remains the dominant mode of operation
- TPPU (Task, Post, Process and Use) is not yet prevalent

Sensor tasking reflects a desire for high resolution to facilitate Detect/ID by human analysts:

- EO is first priority (day clear)
- IR is second priority (night clear)
- SAR (only when another sensor cannot collect, and then SPOT mode only)

Successful development and deployment of ATR would facilitate use of full sensor collection capacity:

- "Assisted" Target Recognition to increase IA efficiency, allowing more high-priority "watch box" tasking

- Fully Automated Target Recognition to exploit lower priority images, enabling full-capacity collection with TPPU doctrine instead of TPED

Findings: Manned Attack Systems

- Manned attack systems present challenges beyond those of ISR systems:
  - Limited time
  - Small displays
  - High stress/workload/multi-tasking
  - Less capable sensors than ISR platforms
  - Legal responsibility for adherence to ROE
  - Potentially severe consequences of error
  - Limited connectivity/bandwidth
- However, for constrained problems, ATR technology can:
  - Assist pilots to find ground targets within small search areas
  - Provide high confidence ID for airborne targets

21

Findings: Autonomous Weapons

- **ATR is not yet ready for fully autonomous weapon release or targeting applications**
  - **ROE and UCMJ considerations are strong constraints on application of ATR to weapons but should not be used as an excuse to stop technology exploration**
  - **High false alarm rates and different consequences of error in different parts of the kill chain should guide near term applications of ATR to weapons**
- **Weapon terminal guidance applications seem well within current ATR technology capabilities and, when used in conjunction with 3-D LADAR sensors, is showing promise**

22

ATR is not yet ready for fully autonomous weapon applications. ATR applications to weapons are technically constrained by FAR that are too high, and by probability of correct ID that is not high enough. In addition, ATR hardware for weapons is physically limited by size, weight, and power. These together lead to constraints on autonomous use imposed by Rules of Engagement (ROE) and by Uniform Code of Military Justice (UCMJ) considerations.

However, weapon applications that use ATR to acquire a target for terminal guidance, either to correct for navigation/target location error or to perform aim point selection, are entirely within the state of the art. Also, when used in conjunction with 3-D LADAR sensors, ATR is showing promise.

Research and technology development will help to reduce these constraints over time. In the interim, the high FAR and different consequences of error in different parts of the kill chain should guide near term applications of ATR to weapons.

Findings: CONOPS & Operations

- **Wide area surveillance is limited by the rate at which the images can be exploited, not by sensor coverage rates**
- **Need to facilitate co-evolution of technology and CONOPS**
  - **CONOPS and ATR performance will interact to set the requirements for Pd and FAR (ROC operating point)**
  - **Some techniques for improving ATR performance have CONOPS implications that must be assessed**
- **Operational ATR programs do not always recognize and fund the associated signature database updates and maintenance**
- **Most ATR technology has been developed to find (and tested against) vehicles and air defense weapons systems – other important targets may require different approaches**

23

As pointed out previously, wide area surveillance area coverage rate is limited by the rate at which the images can be exploited, not by sensor coverage rates.

A likely near-term role for ATR will be to assist people by cueing, but while ATR may increase IA productivity, it will not reduce IA workload.

Investments in analyst support tools (e.g., prioritization of images for exploitation) will have high near term payoff and can lay the foundation for the introduction of ATR.

Some techniques for improving ATR performance have CONOPS implications that must be assessed, and CONOPS and ATR performance will interact to set the requirements for P_d and FAR (ROC operating point). Thus co-evolution of ATR technology and CONOPS would benefit both operators and developers.

Operational ATR programs do not always recognize and fund the associated signature database updates and maintenance. One program updated the sensor between development and production only to discover that the ATR performance with the signatures in the library (collected at great expense during development) was not as good with the production sensor as it had been with the developmental sensor.

Most ATR technology has been developed to find (and tested against) vehicles and air defense weapons systems – other important targets may require different approaches.

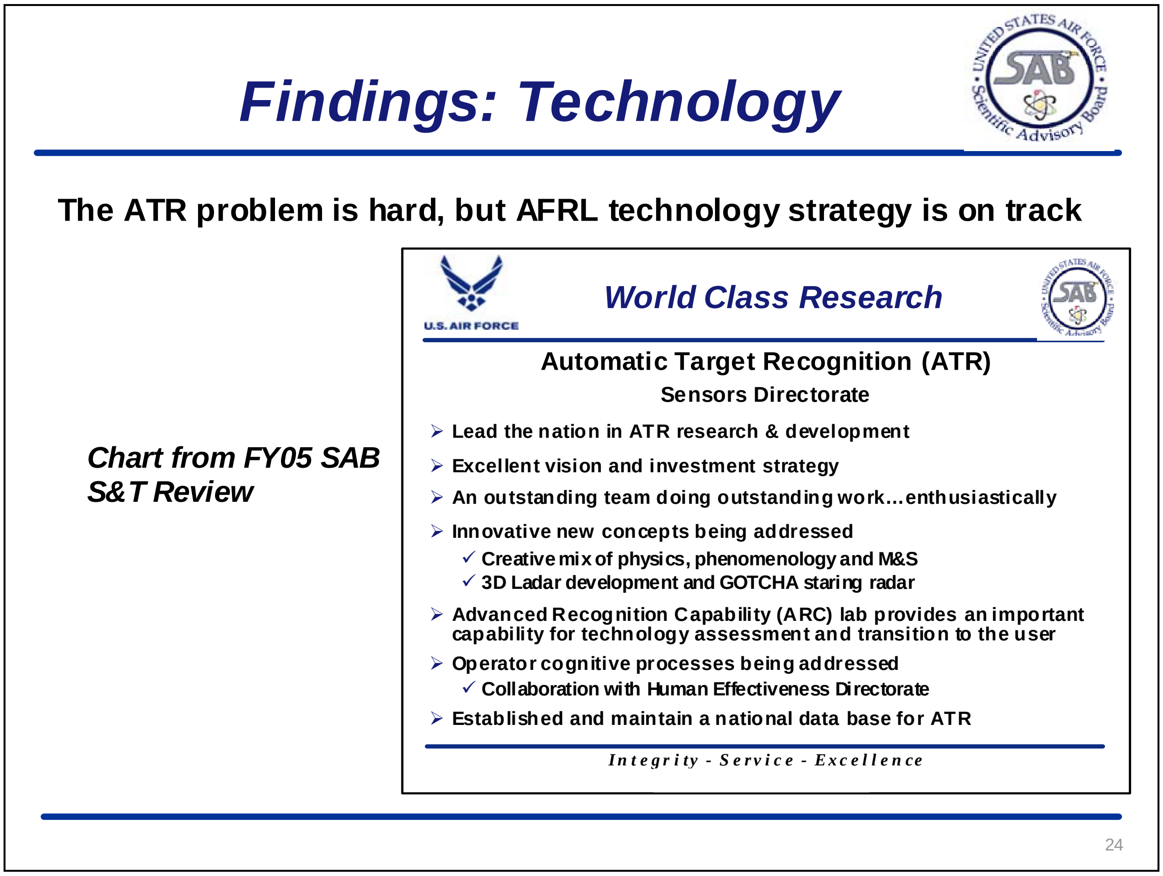

The SAB recently (Oct 04) conducted a S&T Review of the AFRL Sensors Directorate and made special note of the exceptional quality of the ATR work.

The review team found the facilities and people to be top-notch; and the program well structured, based on a clear vision, and strongly coupled to Air Force and outside organizations.

The ATR team successfully transitioned over the past two years from an emphasis on ATR algorithms to an emphasis on systems concepts for enhanced weapon system capabilities. We recommended that this emphasis include use of other sensors and data to aid the recognition decision.

AFRL/SN has approached the ATR problem with a balanced portfolio of physics; phenomenology; and modeling and simulation effort to address the target recognition problem, and two of the more fruitful efforts were their Three Dimensional Laser Detection and Ranging (3D LADAR) development and the staring sensor called the "Gotcha" radar.

The Advanced Recognition Capability (ARC) Lab is a critical node for the assessment of technology developments and transition to the user. In order to address the cognitive operator processes associated with the human-machine issues of ATR, we recommended that they consider netting with the Distributed Mission Operations Center (DMOC).

Section 4: Recommendations

Recommendation #1

- **Assure future ATR maturation efforts are treated in a system context, including (AQ):**
  - **Sensors**
  - **ATR**
  - **Users**
  - **CONOPS**

25

As noted in the findings and illustrated by the Receiver Operating Curves (ROCs), ATR performance is a function of the capabilities of the sensors used, the algorithm by which the sensor outputs are processed (and potentially combined), and how the sensor (collection) and processing sub-systems are employed. In fact, some system concepts (most notably change detection) are enabled or precluded by the user's willingness to employ the system in a particular manner (the collection part of the system in the case of change detection). Similarly, ATR systems built to assist analysts depend on their match to these users' current operating practices or the analysts' willingness to adapt their processes.

Therefore, all the components necessary for realizing ATR capabilities must be brought together when conceiving, developing, and evaluating an ATR system. If any component is neglected, or assumptions made about it without validation, the ATR system's success is put at risk.

Recommendation #2

- **Formalize and implement a mechanism through which the operators and technologists are networked together (on both SIPRNET and JWICS) to enable shared data, cooperative development; and experimental validation of ATR capabilities (XC)**
  - **Include at a minimum:**
    - **DCGS/DGS**
    - **AFRL/Sensors Directorate**
    - **AFRL/Information Directorate**
    - **DMOC**

26

This recommendation can play a big part in accelerating the impact that ATR eventually has on the operator and on overall mission effectiveness. Technologists and developers can use a deeper understanding of the operational perspective to package up the technical capability such that the operator gets more benefit. Operators can use a deeper understanding of what is possible from a technical perspective to create new CONOPS that make better use of the technical capabilities. Co-Evolution of technology and CONOPS can greatly improve the effectiveness of ATR systems. Putting in place mechanisms that promote routine communication between operators and technologists will go a long way towards enabling this co-evolution.

We suggest a minimum of four nodes in this network. Two represent key technologies: the AFRL Sensors Directorate and Information Directorate. Distributed Common Ground/Surface System (DCGS) represents the leading system that can bring these technologies into a system context. And DMOC will provide feedback from the users and also shaping of future CONOPS. This needs to be done a multiple levels of security to insure that we create as many opportunities for insights to occur as possible.

This network, once set up, needs to be utilized for both discovery and development of new capabilities, and also for feedback on existing and developing capabilities.

Recommendation #3

- **Establish an independent activity and process (similar to CIGTF role for inertial navigation systems) to characterize and certify ATR systems (both "automated" and "assisted") prior to acquisition (TE)**

27

When a technology and its implementation become critical to warfighting capability, but before the development engineering community has widely deployed the knowledge, skills and test techniques (and lessons learned), it can be helpful for the Department of Defense (DOD), or a military service to designate a Center of Excellence (COE) for the technology. The COE can be chartered (and funded) to develop and maintain a central capability to test, characterize, and certify systems prior to their acquisition. This was done for inertial guidance systems with the establishment of the CIGTF at Holloman AFB, NM. The objective is to concentrate the relevant expertise and facilities within a single organization, and preferably at a single location, to gain the immediate benefit of common evaluation criteria applied uniformly to acquisition candidates, and to foster productive relationships between the independent test activity and the various development and acquisition organizations that would otherwise develop (or not) duplicate capabilities.

We believe that ATR is a technology that falls into this category, and recommend that the Air Force (as the recognized DOD leader in ATR development) establish an independent activity and process to characterize and certify ATR systems prior to their acquisition.

Recommendation #4

- **Initiate a comprehensive study of ground target identification (the general problem - of which ATR is a small subset) to assess capabilities, progress, and technology needs for this difficult, but important, problem (CC)**

28

The general problem of ground target ID warrants a comprehensive re-examination. While high value targets used to be large military vehicles, high value targets now include pickup trucks and individual people. While the environment of primary interest used to be forests and fields, it now includes cluttered urban areas, with high interest in buried objects and building interiors. Camouflage, concealment, and deception are of increasing concern.

Despite these additional challenges, there is a continual, and perhaps increasing, need for ground target ID.

Further, it is important to understand how current capabilities can be employed to meet this need, and to identify improvements that can be made in the employment of current capabilities for ground target ID. New technologies and technology needs (including but not limited to sensors, platforms, processing, information dissemination, and their integration) should be assessed in the context of ground target ID and its role in the Effects Chain.

Recommendation #5

- **Support ATR-related technology development with continued investment and management attention, to include (AQ):**
  - **Operational and advanced sensor image libraries**
  - **Algorithms, modeling, and simulation-based evaluation**
  - **Independent algorithm and system testing**
  - **Cognitive human-machine interface for users**

29

Significant progress in ATR will require sustained, stable investment rather than our historical practice of episodic investment, usually connected with development of a particular sensor or platform, separated by long "dry spells".

This is particularly true in the collection, calibration, and archiving of the large target and background databases which ATRs need for reference imagery, whether used as direct reference or in the building of analytic target models. This is a major undertaking that must be sustained continuously, rather than being connected to test and evaluation (T&E) of a particular sensor. Databases must be updated as new sensor modalities are developed and targets of interest evolve.

Advances in signal processing bring with them a constant flow of new ATR algorithms and concepts. These can initially be evaluated, and compared, in simulation rather than requiring live implementation.

It is natural for system developers to be overly enthusiastic about ATR performance claims, particularly for new or proposed sensor technologies. It is vital to establish "ground truth" using independent evaluation and testing by objective, experienced technical experts. The AFRL programs are a good step on both counts.

The user interfaces of current ATRs are rudimentary and often substitute one type of confusion for another. This is one reason IAs resist ATR technology. Greater attention is needed to develop advanced user interfaces that are intuitive and compatible with the IAs' operational needs and constraints.

Appendix A: Terms of Reference

Background

The complexity of warfare and the need to shorten the kill chain against difficult targets has increased the need for Automatic Target Recognition (ATR). Applications include both imaging and non-imaging sensors, and range from embedded ATR in weapons, to launch platform target acquisition aids, to ISR exploitation tools, to automated tools which may reduce requirements for IAs. Targets may be fully or partially obscured by competing signals such as foliage, camouflage, or deceptive elements of natural or man-made clutter including electronic countermeasures. The combat ID problem is further increased in complexity in urban environments, by the electronic environment, by target mobility, as well as by weather and the presence of bomb damage debris. Furthermore, no single sensor currently provides a robust solution for all target classes. The ability to rapidly and accurately detect, identify, and track both stationary and mobile targets will enable the Air Force to more effectively conduct operations, reduce sortie rate and minimize collateral damage.

Study Products

Briefing to SAF/OS & AF/CC in October 2005. Publish report in December 2005.

Charter

The study should address the following issues and others it uncovers in the process, and provide appropriate recommendations:

- Assess advances in modeling & simulation capabilities, signature libraries, and exploitation tools
- Assess progress and effectiveness in state-of-the-art S&T in Automatic Target
- Recognition algorithms under different environments
- Assess challenges of ATR and conditions that may limit the effectiveness of ATR
- Assess how far the AF can move from Assisted Target Recognition to ATR
- Define path to accelerate ATR to aid the warfighter

(This Page Intentionally Left Blank)

Appendix B: Wide Area Search Analysis

For this high-level analysis, the search function can be divided into a "find" step and a "confirm" step. The "find" step can cover wide areas with acceptable probability of detection, and a modest false detection rate. The "confirm" step requires more resources, so is applied to double-check all false detections so that only valid detections are passed on.

It is assumed that the density of valid targets is sufficiently small (relative to the rate of false detections) that the false detection rate dominates the workload. It is also assumed that the "confirm" processing perfectly rejects false detections from the "find" processing and consistently accepts detection reports corresponding to actual targets.

Among the system options to be traded off are:

- Multiple platforms: a "find" platform and one or more "confirm" platforms. Here we can trade off false detection rate and coverage rate of the "find" platform, as well as characteristics of the "confirm" platform including number, speed, and sensor range.
- A single platform that changes mode between "find" and "confirm".

The "find" platform is characterized by

α_f (km^2/hr), the area coverage rate of the "find" platform

ϕ (number of false detections/km^2), the false detection density

The average time between false detections is

$$t_d = \frac{1}{\alpha_f \phi} \text{ (hr/false detection)}$$

The average area covered per false alarm is then

$$a_d = \alpha_f t_d = \frac{1}{\phi} \text{ (km}^2\text{/false detection)}$$

Assume that the average distance between false detection locations is

$$s_d = \sqrt{a_d} = \frac{1}{\sqrt{\phi}} \text{ (km/false detection)}$$

Assume that the "confirm" sensor has a range of r_c (km), so the average distance that the "confirm" platform needs to travel between false detection locations is

$$s_d - 2r_c \text{ (km/false detection)}$$

The speed of the "confirm" platform needed to keep up with false detections is

$$\bar{v}_c = \frac{s_d - 2r_c}{t_d} = \left(\frac{1}{\sqrt{\phi}} - 2r_c\right)\alpha_f \phi = \alpha_f \sqrt{\phi} - 2r_c \alpha_f \phi \text{ (km/hr)}$$

If the "confirm" sensor range is zero (i.e., the sensor must be overhead the point it is viewing), the speed of the "confirm" platform needed to keep up with false detections is

$$\bar{v}_c = \alpha_f \sqrt{\phi} \text{ (km/hr)}$$

The actual speed of the "confirm" platform is

$$v_c \text{ (km/hr)}$$

If it is no use doing "find" in an area without "confirm", and no use doing "confirm" in an area without "find", then the effective rate of doing both find and confirm is

$$\alpha_{f+c} = \min\left(1, \frac{v_c}{\bar{v}_c}\right)\alpha_f \text{ (km}^2\text{/hr)}$$

Presumably, if there are n_c "confirm" platforms, the effective rate of doing both find and confirm is

$$\alpha_{f+c} = \min\left(1, \frac{n_c v_c}{\bar{v}_c}\right)\alpha_f \text{ (km}^2\text{/hr)}$$

Now, consider an alternative situation where the same platform performs both "find" and "confirm", but cannot perform "find" while it is performing "confirm".

The time taken for "confirm" is

$$T_c \text{ (hr)},$$

so the effective rate of doing both find and confirm is

$$\alpha_{f+c} = \alpha_f \frac{t_d}{t_d + T_c} = \alpha_f \frac{1}{1 + \alpha_f \phi T_c} \text{ (km}^2\text{/hr)}$$

Numerical Examples

Following are some numerical examples based on $\alpha_f = 4000\ (\text{km}^2/\text{hr})$. We'll probably want to adjust some of the values, but the plots provide some insight and the combined set of numerical values appears to be in the ballpark of interest. The first plot is for a single "confirm" sensor, with different speeds for the "confirm" platform, when the "confirm" sensor has zero range (so it has to be directly over the location being viewed).

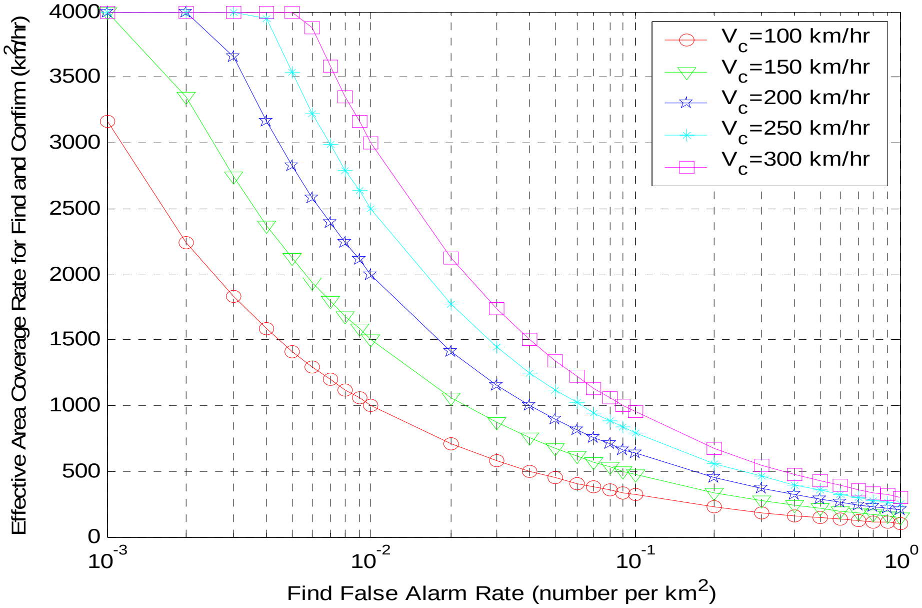

The second plot is for the same sensor switching between "find" and "confirm" modes, with different amounts of time consumed in "confirm".

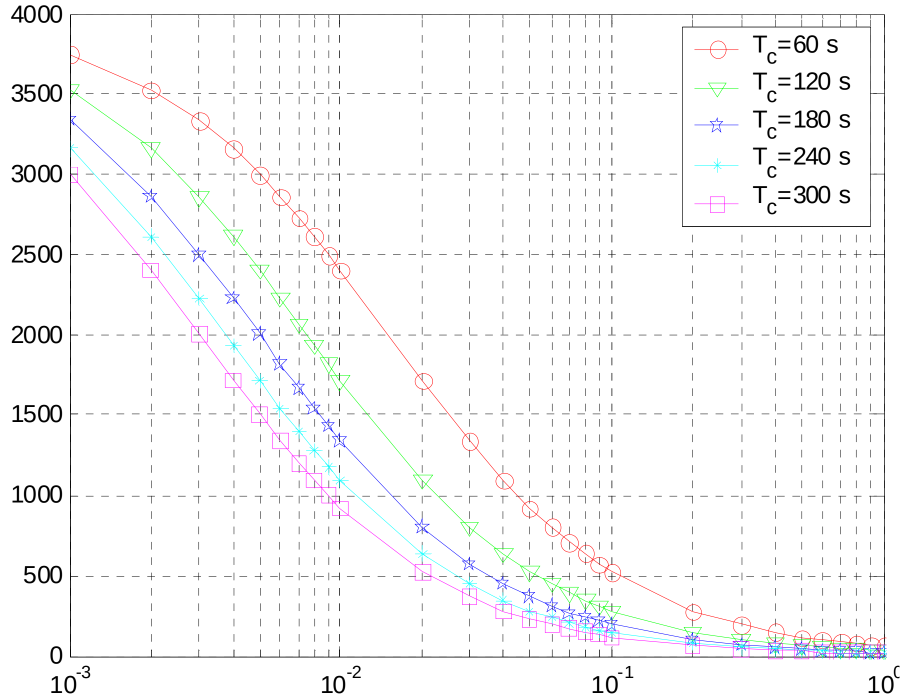

Appendix C: SAR ATR State of the Art

Air Force and other investments in ATR for Synthetic Aperture Radar (SAR) Imagery have produced algorithms for multiple sensors and multiple applications. The maturity of these algorithms is described in terms of the algorithms and their performance.

Algorithms

Industry and government researchers have developed automated or semi-automated approaches to detecting and identifying targets. These algorithms were intended for very different applications, from prioritizing imagery for humans in ground stations to cueing aircrew to likely targets, to autonomous combat ID. Despite these intended differences in functionality and application, there are similarities in the development processes and the algorithms resulting from them.

In each case, there is a design process that selects the form of target knowledge (features) the algorithm will use to discriminate among targets and between targets and non-targets. There is an off-line process (often called training) that develops the target knowledge representation and establishes the parameter values (thresholds) used for discrimination. Finally, there is an on-line process (the algorithm) to compare extracted information from incoming data with that stored in the off-line process, according to the thresholds set in the training process. Differences features and in the training approach are often greater than the differences in the algorithms themselves.

Feature varieties include parametric features, templates, and models. Parametric features are physically derived (e.g. target dimensions) or abstract descriptors that developers believe will discriminate between targets. These features are compact, but their extraction in the on-line process is challenges, and they may suffer from both fragility and limited discrimination power. Still, their simplicity makes them suitable for initial screening. Templates are instances of target signatures used to compare with incoming imagery. Sets of templates across a variety of anticipated conditions (e.g. aspect angle) compose a signature database, which may be derived from multiple measured and synthetic imagery sources. Template features are the most commonly employed. Models are geometric target models, along with signature prediction capability, that allows the on-line algorithm to generate signatures on the fly and hone in on the correct target type.

The training procedure used in current algorithms is generally not an automated or well-understood process. Developers employ their experience as well as experimentation to derive the thresholds set within an algorithm to determine “target likeness” and discard non-targets. Technical progress in this area would allow more automated signature database updates and more accurate performance prediction.

The on-line algorithm, then, compares features observed in incoming data with those stored in training to determine target presence and type. Most algorithms perform an initial detection process to limit the image area investigated by the more computationally intensive ID process.

The out put of the initial process typically includes many false alarms, so even an algorithm intended for detection or cueing will typically include an ID function to limit false alarms.

Performance

Performance measurement keys for ATR are the metrics used and the conditions over which these metrics are understood. Metrics commonly employed include detection and FAR, used to capture the algorithm's ability to find targets, and ID and false target probabilities, used to capture the algorithms ability to further classify or discard found targets.

An algorithm's detection and false alarm rate are related, and neither reported alone is meaningful. Most algorithms allow tuning of this "operating point," and the receiver operating characteristic (ROC) reports all possible operating points for an algorithm. A critical part of understanding an algorithm's ability to find targets is the definition of targets. Separation of large vehicles from natural clutter is much easier than separating a specific set of large vehicles from other vehicles as well as clutter.

An algorithms ID performance is characterized by its ability to correctly sort targets by type (ID probability) and to reject non-target objects (false target probability). The ability to reject non-target objects is critical to ATR employment; without it, civilian, industrial, and allied vehicles will be called targets by the ATR. Most ATR development and testing fails to explore this critical issue.

An algorithm's performance, as measured by these metrics, will vary according to sensing, environmental, and target conditions. In addition to the obvious challenges presented by certain conditions – camouflage, for example – the algorithms may face more subtle challenges under seemingly benign conditions not represented in training data.

For synthetic aperture radar, sensing conditions include depression and squint angles, sensor anomalies and noise level, focusing artifacts, and compression effects. While the application constrains some of these parameters, others will not match training data. Environment conditions include background contents, terrain type and relief, and intentional deception. Again, application constrains some of these factors, but not all background combinations will be observed in training data. Finally, the adversary defines the target conditions, including the target types and variants and their deployment along with aging, damage, articulation, and camouflage of the target. It is important to understand performance in the context of the breadth and difficulty of these conditions as well as the relationship between test conditions and training conditions. Ultimately, the relationship between operational conditions – anticipated or not – and available training data may drive performance.

Algorithm performance has been carefully measured for 0.3 m resolution sensing. Third party evaluation of multiple developers' algorithms on sequestered data sets yields consistent results for both finding and identifying targets. At a high target detection rate (0.9), algorithms generate about one false alarm per square kilometer resulting from clutter. At a reduced target detection rate (0.5), this FAR is reduced by nearly two orders of magnitude. The ROC curve showing other possible operating points is shown below, along with expected detection performance at 1.0m resolution.

Conditions for 0.3m Resolution Results	
12 GOB Targets	
15 confuser (non-target) vehicles	
Sensing Conditions	10, 15, 25° depression angles
Environment Conditions	Multiple Backgrounds, Obscuration
Target Conditions	Multiple Serial Numbers
	Version Variants and Configurations
	Door, Turret and Hatch Articulation

ID Performance	
@P_d=0.9	@P_d=0.5
FAR=1/km^2	FAR=0.02/km^2
P_{ID}=0.85	P_{ID}=0.95
P_{FA}=0.5	P_{FA}=0.1

The algorithm's ID performance can also be observed at each point on the ROC curve. At high detection rates, current algorithms can correctly sort targets by type 85% of the time. At reduced detection rates, where only the "easier targets" must be classified, the ID rate rises to as much as 95%. Troubling, however, is the algorithm's limited ability to reject non-targets (which could be allies, industrial vehicles, or civilians). At high target detection rates, 50% of non-targets presented to the algorithm are ruled targets by the algorithm, and labeled as a type present in training. When only target sized vehicles are considered, this rate rises to 90%! This result warns the community about autonomous use of current algorithms, particularly in areas with civilian or friendly activity.

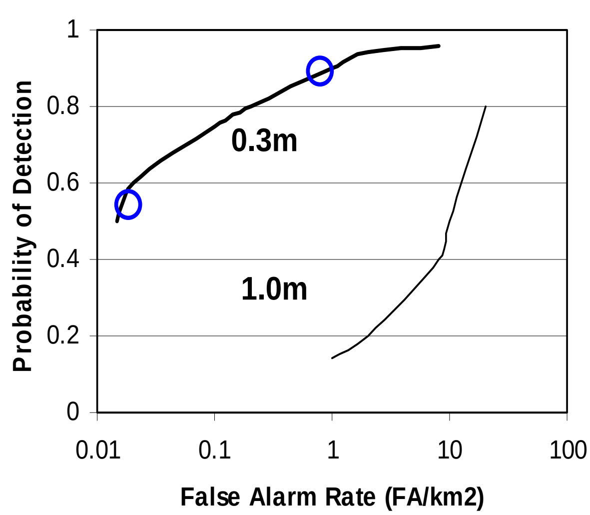

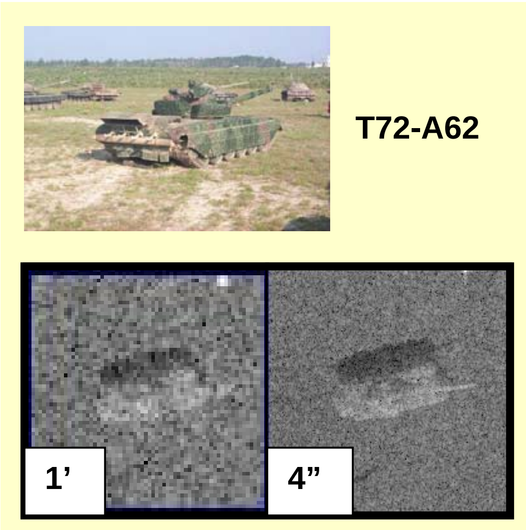

Initial experimentation on higher resolution imagery does show promise for reducing both clutter and non-target false alarms and increasing ID rate. Still, providing training examples of non-target vehicles from measurements or synthetic data is the surest way to reduce the occurrence of non-target false alarms.

Near term employment of SAR ATR will require constraining the application by limiting the target set, training on non-targets, operating at reduced detection rates, and ensuring human interaction with the semi-automated system.

Appendix D: Multisensor ATR

Several approaches are being considered for improving ATR (detection, tracking, ID) by bringing together multiple looks, modes or sensors. These approaches promise performance benefits but warn of supporting technology and CONOPS that would be required to enable them. This appendix discusses potential challenges as well as performance potential for *alert/confirm cueing* and *multisensor reasoning*. This technique places minimal requirements on supporting technology while providing highly reliable cues.

Alert/Confirm Cueing

In alert/confirm cueing, cues from a sensor with large area coverage, but coarse resolution, are interrogated by narrower field of view, but high-resolution sensors to identify targets. This concept can be applied to multiple resolutions of sensing from the same platform, to different sensors on the same platform, or to separate platforms. Since the information passed from one sensor to the other is essentially location, bandwidth requirements are low. The location accuracy with which the alert sensor must "steer" the confirm sensor is coarse, only requiring that the confirm sensor can reacquire the nominated location. Finally, understanding the performance of the component processes (detection and false alarm performance of the alert sensor and its processing, and ID and confuser rejection of the confirm sensor and its processing) provides accurate prediction of the system's performance.

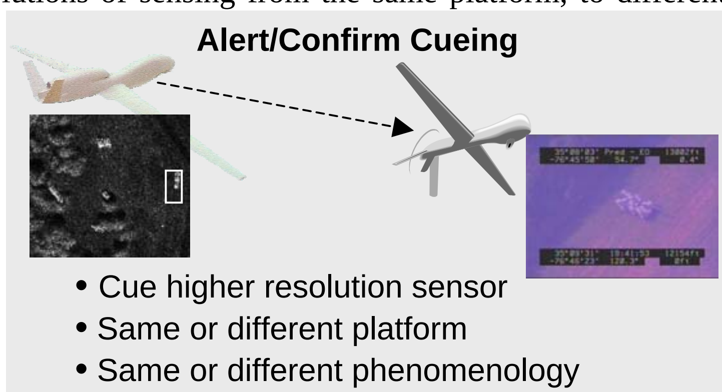

Multisensor Reasoning

In multisensor reasoning, imagery collected from multiple sources may provide rich target information. Ground phenomena that cause false alarms differ across sensors, so multisensor reasoning may reduce false alarms. For targets, sources may provide detail on different target characteristics, including geometry, surface properties, and internal characteristics. This additional detail may allow higher ID performance as well as additional target characterization. However, significant challenges in registration, signatures,

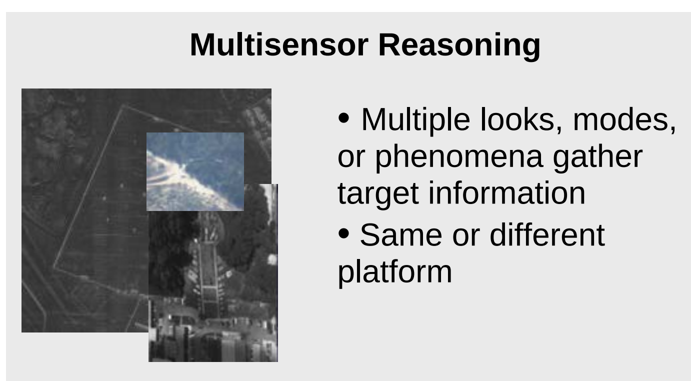

evidence accrual, and communications will make automated exploitation difficult. Performance benefits are not currently quantified.

Appendix E: Creating the Receiver Operating Characteristic (ROC) Curve

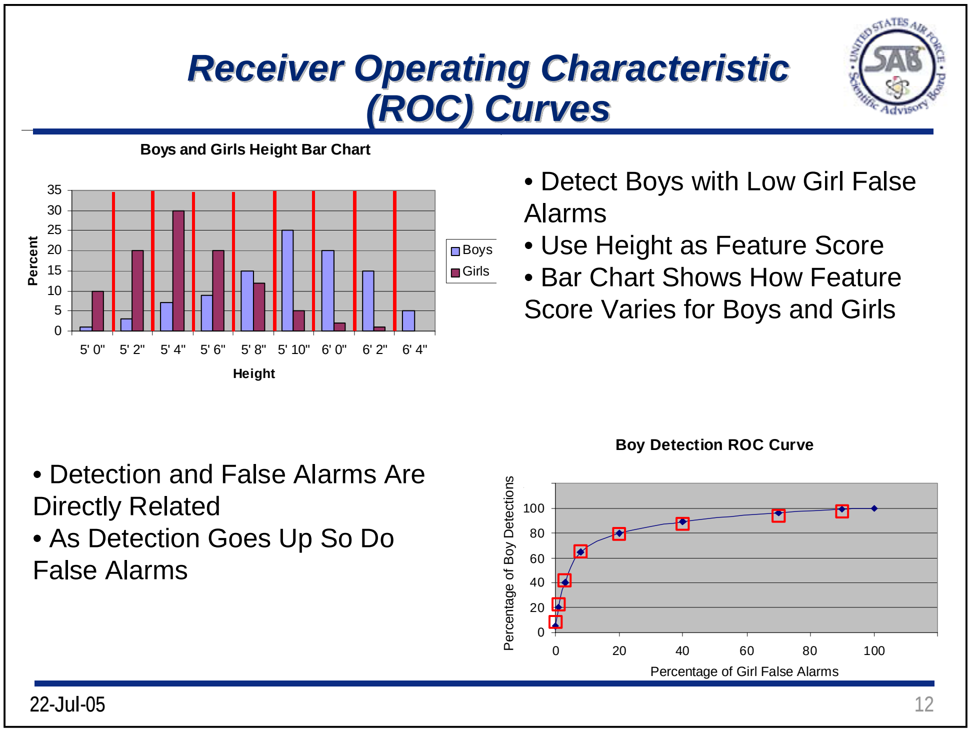

In order to understand a Receiver Operating Characteristic (ROC) curve, it is useful to generate one using an example. Let's use height to discriminate boys from girls. First, we need to understand how boys and girls are distributed according to height. This distribution is shown in the bar chart in the upper left hand corner entitled "Boys and Girls Height Bar Chart". Note that the total population of girls shown in red adds to 100%, as does the total population of boys shown in blue. In technical terms, these bar charts represent the probability density functions of the boys and girls.

Now, in order to generate a ROC curve we start with the bar chart and set a threshold. The red vertical lines on the bar chart illustrate these thresholds. So if we set a threshold at 6'3" (we ask what percentage of boys and girls are over 6'3"), we find that 0% girls are over 6'3", and 5% of the boys are over 6'3". This point (0% girls, 5% boys) is plotted on the curve in the lower right hand corner. It is the farthest point to the left on the chart entitled "Boy Detection ROC Curve". It means that if we decide to set the threshold at 6'3", we will detect 5% of all the boys with no girls detected or no girl false alarms.

Now, if we want to detect more boys, we can set the threshold at 6'1" (which is the second vertical line from the right on the upper left chart entitled "Boys and Girls Height Bar Chart"). Here we can see that there are now 20% of the boys taller than the 6'1" threshold and only 1% of the girls. If we plot this point on the "Boy Detection ROC Curve", we obtain the second point from the left (1% girls, 20% boys) which indicates that a threshold of 6'1" will detect 20% of the boys but will false alarm on 1% of the girls.

The lower we draw our threshold (the farther we move the red line on the "Boys and Girls Height Chart" to the left), the more boys we detect, but also we obtain more girl false alarms. Hence, if we choose the red line furthest to the left, we detect 99% of the boys but false alarm on 90% of the girls. Clearly, each red line or threshold on the "Boys and Girls Height Bar Chart" corresponds to a red box on the "Boy Detection ROC Curve". Also, it is clear that detection and false alarms are directly related. As the detection percentage increases so does the probability of false alarm.

The ROC curve is an important way of measuring the performance of an ATR system that performs a detection function. Detection and false alarms are directly related. Any performance claim that gives one without the other is meaningless. Further, it is also important that the conditions of the test be described as thoroughly as possible. For example, if we knew that the conditions of this test were that the boys were from the boys' basketball team and the girls were from the girls' gymnastic team, we would expect a different distribution of heights on the bar chart and a more favorable ROC curve (i.e., more boy detections, less girl false alarms for each threshold setting).

Appendix F: Definitions

3D LADAR (Three Dimensional Laser Detection and Ranging) imaging sensor enables target search, location, and classification amidst obstructions and clutter. The 3D LADAR accomplishes this by not only rendering an object's picture, but also its three-dimensional shape and relative distance from the observer.

Apache Longbow (AH-64) is the Army's heavy division/corps attack helicopter. The weapon system incorporates a millimeter wave fire control radar (FCR), radar frequency interferometer (RFI), fire-and-forget radar-guided HELLFIRE missile and cockpit management and digitization enhancements. Apache Longbow conducts rear, close, and deep precision strike to include distributed operations, precision strikes against relocatable targets, and provides armed reconnaissance when required in day, night, obscured battlefield and adverse weather conditions.

ATR (Automatic Target Recognition), n., A process whereby sensor information is converted into the detection and categorization of targets of military interest. This applies to both Autonomous (without benefit of human intervention) and Aided (human in the loop) Target Recognition. (Automatic Target Recognizers Working Group - 1980's)

Change detection compares the incoming sensor data to pre-stored background data to detect changes (e.g., targets moving in or out of an area). For change detection, the background database requires at least one previous image of the area.

False Alarm is 1 of 3 elements that characterize detection performance (the other 2 elements are correct detections and missed detections). A false alarm is the result of the system incorrectly identifying a wrong target as a mission target.

Global Hawk is an Unmanned Aerial Vehicle (UAV) providing near-real-time, high-resolution, intelligence, surveillance and reconnaissance imagery. Cruising at extremely high altitudes, Global Hawk can survey large geographic areas to give military decision-makers information about enemy location, resources and personnel. Once mission parameters are programmed into Global Hawk, the UAV can autonomously taxi, take off, fly, remain on station capturing imagery, return and land.

Gotcha Radar is a persistent radar with temporal constancy and angular diversity. Gotcha Radar can be used by Global Hawk, providing a continual scan of a particular area of interest.

ID (Identification) "...validating the threats as worthy of continued attention.... establish and declare confidence that the threat is real and poses potential danger to our forces." (AOC CONOPS Mar 01, Paragraph 3.2.5.2.1.2, definition of Fix)

Receiver Operating Characteristic (ROC) curve describes the range of performance options that are available from a particular ATR system. The developer (and in well designed systems, the user) selects the operating point.

Scene Matching Area Correlator (SMAC/DSMAC), the optomechanically implemented SMAC, and its digital grandchild DSMAC, originated at the Naval Avionics Facility Indianapolis in the late 1960's and was fielded in Tomahawk beginning in the 1980's.

Terrain Contour Matching (TERCOM) is an ATR technique that uses the uniqueness of terrain altitude profiles to update cruise missile navigation systems based on radar altimeter measurements. It was developed beginning in the mid 1960's and was first fielded in the Air Launched Cruise Missile (ALCM) in the 1980's.

Appendix G: Acronyms and Abbreviations

3D LADAR	Three Dimensional Laser Detection and Ranging
AF	U.S. Air Force
AF/CC	Chief of Staff of the Air Force
AFRL	Air Force Research Lab
AFRL/SN	AFRL Sensors Directorate
AFSAB	Air Force Scientific Advisory Board
ALCM	Air Launched Cruise Missile
APC	Armored Personnel Carrier
AQ	Secretary of the Air Force, Acquisition
ARC	Advanced Recognition Capability
ATR	Automatic Target Recognition
BDA	Battle Damage Assessment
C4	Command, Control, Communications, and Computers
CC	Air Force Chief of Staff
CC&D	Countermeasures, concealment, and deception
CIGTF	Central Inertial Guidance Test Facility, Holloman AFB, NM
COE	Center of Excellence
CONOPS	Concept of Operations
DCGS/DGS	Distributed Common Ground/Surface System/Distributed Ground System
DMOC	Distributed Mission Operations Center
DOD	Department of Defense
DSMAC	Digital Scene Matching Area Correlator
EO	Electro-optical
FAR	False Alarm Rates
FFRDC	Federally funded research and development center
IA	Image Analyst
ID	Identification
INS	Inertial Navigation System
IR	Infrared

ISR	Intelligence, Surveillance, Reconnaissance
JWICS	Joint Worldwide Intelligence Communications System
P_d	Probability detection
ROC	Receiver Operating Curve
ROE	Rules of Engagement
S&T	Science and Technology
SAB	Air Force Scientific Advisory Board
SAF/OS	Office of the Secretary of the Air Force
SAR	Synthetic Aperture Radar
SIPRNET	SECRET Internet Protocol Router Network
SMAC	Scene Matching Area Correlator
T&E	Test and evaluation
TERCOM	Terrain Contour Matching
TPED	Task, Process, Exploit, Disseminate
TPPU	Task, Post, Process, Use
UCMJ	Uniform Code of Military Justice
XC	Secretary of the Air Force, Warfighting Integration and Chief Information Officer

Appendix H: Initial Distribution

Air Force Leadership

Secretary of the Air Force
Chief of Staff of the Air Force
Under Secretary of the Air Force
Vice Chief of Staff of the Air Force

Air Force Secretariat

Secretary of the Air Force, Acquisition (SAF/AQ)
Secretary of the Air Force, Warfighting and Chief Information Officer (SAF/XC)

- Air Force C2ISR Center

Air Staff

Assistant Vice Chief of Staff of the Air Force
Director of the Air National Guard
Chief of Air Force Reserve
Scientific Advisory Board Military Director
Chief Scientist of the Air Force
Deputy Chief of Staff of the Air Force Air and Space Operations

- ISR Directorate
- Operations and Training Directorate
- Requirements Directorate

Air Force Major Commands

Air Combat Command
Air Education & Training Command
Air Force Materiel Command
Air Force Space Command
Air Force Special Ops Command
Air Mobility Command
Pacific Air Forces
U.S. Air Forces in Europe
Air Force Reserve Command

Other Air Force Elements

Air Force Research Laboratories

- Information Directorate
- Munitions Directorate

- Sensors Directorate

480th Intelligence Wing

Executive Office of the President

National Security Council

Office of the Secretary of Defense

Under Secretary of Defense (Intelligence)

Joint Chiefs of Staff

Chair, Joint Chiefs of Staff
Vice Chair, Joint Chiefs of Staff
Joint Chiefs of Staff, Director of Intelligence
Joint Chiefs of Staff, Director of C4 Systems
Joint Chiefs of Staff, Director of Operational Plans and Interoperability

Defense Agencies

Defense Information Systems Agency
Defense Advanced Research Projects Agency

Combatant and Regional Commands

U.S. Central Command
U.S. European Command
U.S. Joint Forces Command
U.S. Northern Command
U.S. Pacific Command
U.S. Southern Command
U.S. Special Operations Command
U.S. Strategic Command
U.S. Transportation Command

Intelligence Community

Air Force Intelligence Agency
Central Intelligence Agency
Defense Intelligence Agency
Department of State – Bureau of Intelligence and Research
Marine Corps Intelligence Activity
National Geospatial-Intelligence Agency

National Reconnaissance Office
National Security Agency
Office of Naval Intelligence

Advisory Boards
Army Science Board
Defense Policy Board
Defense Science Board
Naval Research and Advisory Committee
Naval Studies Board

Industry
SAIC

- Technology Research and Integration Business Unit

BAE Systems
Johns Hopkins University Applied Physics Lab (JHU/APL)
Northrop Grumman

- MS Standards Lab
- Radar Systems Architects

MITRE

- Center for Integrated Intelligence Systems

Scitor Corporation

- National Imagery Systems Sector

(This Page Intentionally Left Blank)

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and manipulating the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (*Leave Blank*)	2. REPORT DATE 1 July 2005	3. REPORT TYPE AND DATES COVERED Final, January 2004 – July 2005		
4. TITLE AND SUBTITLE ***Automatic Target Recognition, Executive Summary and Annotated Brief (PR)***			5. FUNDING NUMBERS	
6. AUTHOR(S) Dr. Llewellyn Dougherty				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) HQ USAF/SB 1180 AF PENTAGON RM 5D982 WASHINGTON, DC 20330-1180			8. PERFORMING ORGANIZATION REPORT NUMBER SAB-TR-05-02-PR	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) SAF/OS, AF/CC AIR FORCE PENTAGON WASHINGTON, DC 20330-1670			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Public Release			12b. DISTRIBUTION CODE A	

ABSTRACT (Maximum 200 Words)

Automatic Target Recognition. The complexity of warfare and the requirement to shorten the kill chain against difficult targets has increased the need for Automatic Target Recognition (ATR). No single sensor currently provides a robust enough capability of detecting all classes of targets through the camouflage and deceptive elements of natural and/or man-made clutter. ATR is a sub-element of a larger class of problems known collectively as "Target Identification". This study concentrated on image-based techniques as an ATR sub-discipline of target ID.

The study panel concluded that no "silver bullet" exists that will dramatically improve ATR performance in the near term. However, more work and understanding in very specific areas could improve the way the Air Force uses ATR today. For example, CONOPS and technology can work synergistically (e.g. change detection) to improve mission capabilities; ATR can be an aid to, but not a substitute for, human judgment by both pilots and image analysts; and ATR is essential to better match image exploitation capabilities with collection capacity.

14. SUBJECT TERMS Automatic Target Recognition, ATR, Sensor, SAR, Synthetic Aperature Radar, Change Detection, Multi-Sensor, Wide Area Search, ISR, Persistence, Attack			15. NUMBER OF PAGES 62
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT PUBLIC RELEASE