AATR for CT Based EDS

David Castañón

Boston University

T. Breckon
Q. Wang
K. Ismail
U. Durham (UK)

D. H. Ye
C. Bouman
Purdue 1

D. Paglieroni
H. Chandrasekaran
C. Pechard
H. Martz Jr.
LLNL

A. Kac
A. Manerikar
Purdue 2

This material is based upon work supported by the U.S. Department of Homeland Security, Science and Technology Directorate, Office of University Programs, under Grant Award 2013-ST-061-ED0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security

So What, Who Cares?

BOSTON UNIVERSITY

- Desire to adapt EDS system performance to changes in threat requirements
  - New threats, new Regions of Responsibility for existing threats, new performance requirements
  - Would like this adaptation to happen rapidly (order of days, not months).
- ALERT conducted a task order experiment in Adaptive Automated Target Recognition, sponsored by Drs. Laura Parker and John Fortune, DHS
  - Five performers: Boston University, Livermore Laboratory, University of Durham (UK), and two teams from Purdue University
  - Each team developed different approaches for ATR and adaptation to changes in performance requirements
  - Approaches were evaluated using both training and sequestered unclassified data from medical scanner
- Approaches had varying degrees of success
  - Identified issues in adaptation, validation, potential certification of approaches
  - Lessons could be valuable for evolving existing EDS systems to increase adaptability

The Data

- 3-D volumetric images of containers, obtained by IMATRON Scanner: single spectrum
- Objects of interest in training data
  - Rubber sheets of different thickness, widths
  - Bulk saline bags, different sizes
  - Bulk clay, different sizes
- Training data is available, distributed by ALERT to promote further research
  - Sequestered data for evaluation not distributed
- Additional objects of interest
  - Several compounds of different sizes, designed by Livermore Laboratory
  - Included only in sequestered data

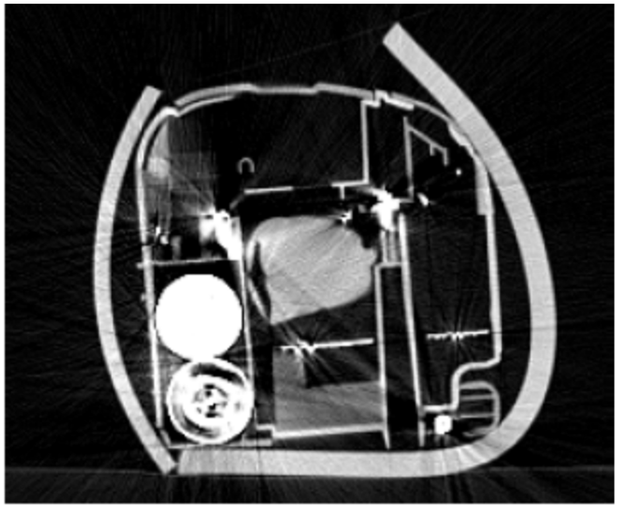

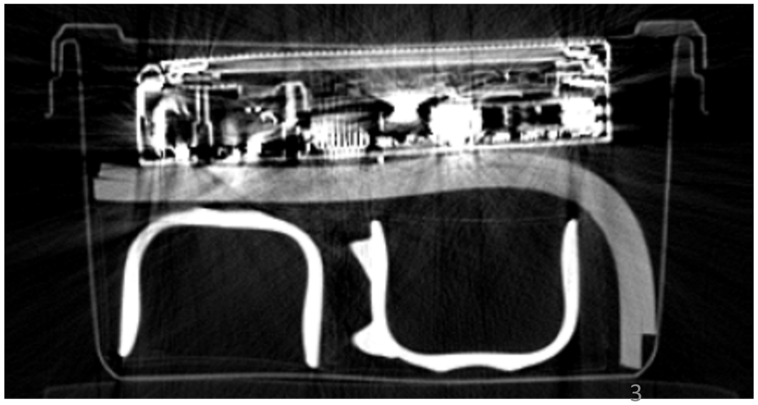

Adaptive ATR Tasks

- Requirements: adapt EDS ATR system to changes in:
  - Desired P_D, P_{FA}
  - Definition of threats: mass, density, thickness, type, …
  - Differential P_D, P_{FA} per threat class
  - New threats from specifications without training data
- Text specifications provided to AATR
  - AATR required to modify ATR to meet specifications
  - Able to use training data to cross-validate predicted performance
  - Limited significantly in predicting performance for objects with no training data

```
[ORSName] ors3.txt
[ORSVersion]  $Revision: 1.4 $
[ORSSynopsis] PD_Weight for clay = 0.25
[NOOI] 2
[PFA] 0.1

# OOI 1: saline
[Synopsis] saline
[Type] t
[SubType] saline
[SubTypeCode] s
[Form] N/A
[MassMin] 137
[MassIncrement] 1000
[RhoMin] 1050
[RhoMax] 1215
[ThicknessMin] 6
[ThicknessIncrement] 10000
[Texture] homogeneous
[PD] 0.9
[OOIMerge] no

# OOI 2: clay
[Synopsis] clay
[Type] t
[SubType] clay
[SubTypeCode] c
[Form] N/A
[MassMin] 80
[MassIncrement] 1000
[RhoMin] 1530
[RhoMax] 1715
[ThicknessMin] 6
[ThicknessIncrement] 10000
[Texture] homogeneous
[PD] 0.2
[OOIMerge] no
```

ATR Approaches

- Typical of ATR processes, includes segmentation, feature extraction, and classification
- Each team developed independent approaches to these functions
  - Different segmentation, features, classifiers
- Classification approaches:
  - One vs. all SVM yielding probability of type, used for classification
  - Gaussian sum voxel classification followed by consensus smoothing of likelihoods, then segmentation and eventual classification
  - K-nearest neighbor classifiers
  - Random Forests

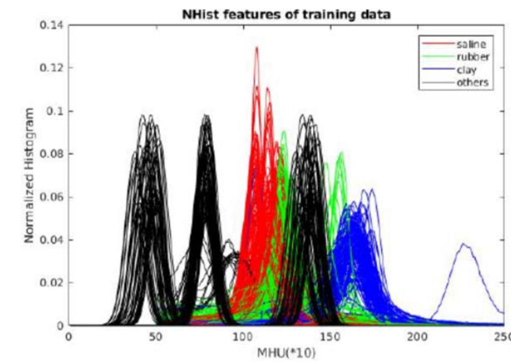

Histogram of intensities/class

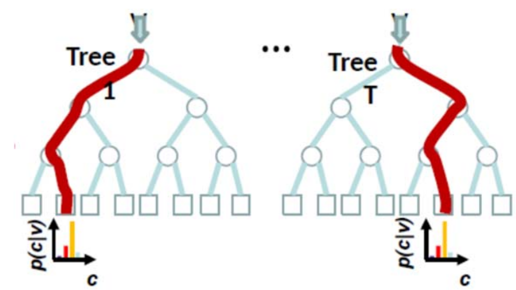

Random Forest Classifier

Evaluation Metrics

BOSTON UNIVERSITY

- Output of ATR: Volumetric image with location of detected threat volumes in bag
- Ground truth used in scoring: Volumetric image of true threat volumes in bag, hand developed using videos of packing container plus manual recognition
- Fundamental metrics:
  - Detection: Significant overlap between detected threat volume and ground truth threat volume
  - False alarm: Reported threat volume that does not have corresponding ground truth threat volume

Ground Truth

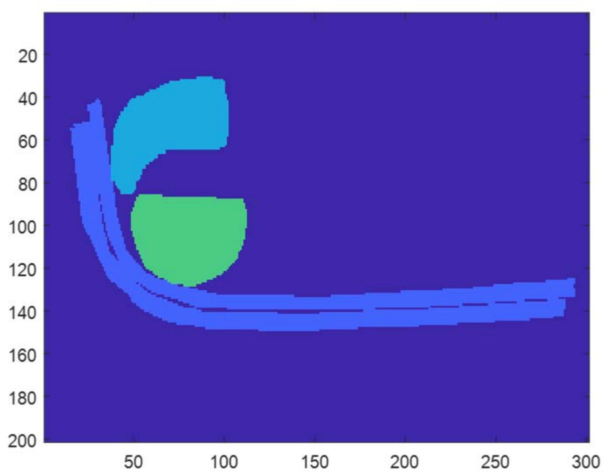

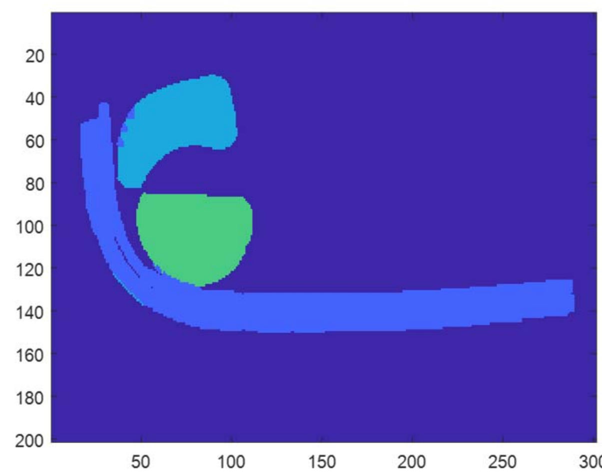

ATR Output

Adaptation to Requirements for Known Threats

BOSTON UNIVERSITY

- Problem: Changing Prob. Detection/False alarm tradeoffs for threats in training data
  - Increase importance of some threats, decrease others
  - Change region of responsibility (minimum mass, thickness, density spread, …)
- Approaches:
  - Change thresholds for decisions, keeping similar processing of features, structures (Durham, Purdue 1, LLNL)
  - Reweight training data, retrain classifiers while cross-validating performance
    (Purdue 2, BU)

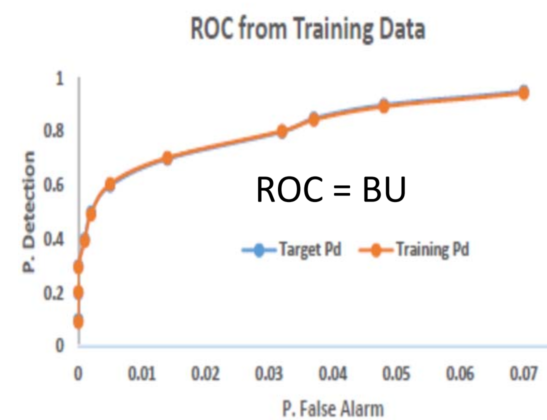

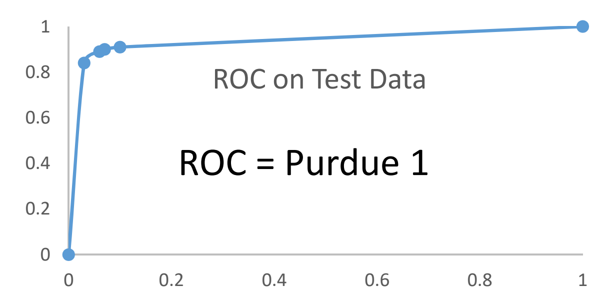

Adaptation to Requirements for Unknown Threats

BOSTON UNIVERSITY

- Unknown threats: Not present in training data
  - Only information is on simple RoR parameters:
    - Density range, Mass range, Thickness range
- Very different approaches investigated:
  - Randomly generate new training data using RoR guidance → integrate with existing training data (Durham, Purdue 2)
  - Randomly generate parameters of classifiers (e.g. Gaussian sum parameters, Feature samples for k-nearest neighbors) → generate single merged classifier (Purdue 1, LLNL)
  - Design separate classifier using reduced feature set, integrate into overall structure using parallel paths (BU)
- Key issue: don't have data to cross-validate performance!
  - Can assess P_{FA} because of available background training data, but not P_D
  - Addressed in this effort by tuning using multiple attempts, but not practical in real scenarios with unknown threats

Performance of Approaches for Unknown Threats

BOSTON UNIVERSITY

- Evaluated with only threat the unknown threats
- Some approaches based on simulating data were very sensitive to assumptions not provided in requirements
- Performance of separate classifier (team 5) with reduced feature set as good or better than alternatives.

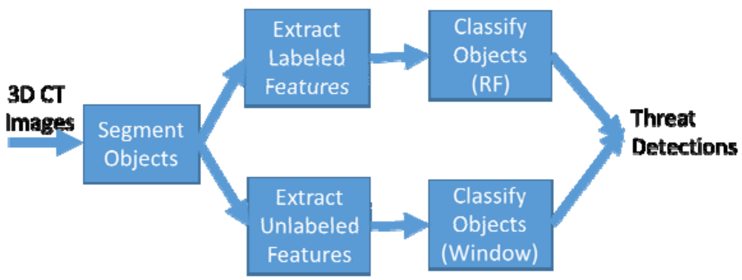

					Durham		Purdue 1		LLNL		Purdue 2		BU	
	Density Range	Minimum	Req.	Req.	Team 1		Team 2		Team 3		Team 4		Team 5	
OOI	(MHU)	Mass (g)	PD (%)	PFA (%)	PD	PFA	PD	PFA	PD	PFA	PD	PFA	PD	PFA
A1	380-525	42	90	10	76	12	83	14	94	11	26	13	89	1
A2	770-810	67	90	10	100	46	100	13	85	4	71	47	100	5
A3	1300-1375	174	90	10	92	15	100	12	96	2	28	38	92	11
A4	1350-1430	183	90	10	100	11	100	6	80	1	25	70	100	0

Ongoing near-term work

- Testing on new materials with requirements for mixture of known/unknown threats
- Evaluating utility of providing additional information on known threats
  - Small examples from sequestered data
  - Additional description of RoR
- Report documenting effort available (any time now) from ALERT on request

So what was learned?

- Can extend ATR algorithms to modify performance requirements
  - Does require access to training data to tune/evaluate adaptation
- Key issue: How to validate/verify detection performance against new threats with limited specifications
  - Cannot do iterated performance testing with sequestered data
  - Consider use of simulated data embedded in stream of commerce to generate cross-validation
  - Can also consider generation of simulated image data for training → not explored in current effort that used only reconstructed images.
- Use of separate classifier for new threats provides easy path for expansion of existing EDS systems
  - Certified EDS component not modified
  - Can provide interim capability while additional threat characterization is obtained

Questions from Carl

BOSTON UNIVERSITY

- What were the comments from TSA and the vendors at the program review?
  - Wished performers used common early processing to show differences in AATR
  - Thought approaches were reasonable, similar to how rapid response would be implemented in own EDS systems
- What will it take to deploy an AATR?
  - Trust: demonstration that performance can be achieved when changing software parameters to respond to new requirements
- What changes will TSA have to make to deploy an AATR?
  - Develop approach to verify/certify adaptive algorithms
  - Similar to certifying feedback control systems, etc – much recent work in area
- How does ALERT's AATR compare to other methods for dealing with an adaptive adversary?
  - Focus on rapid exploitation of intelligence data to close vulnerability.
  - Don't know too much about other methods…