AATR for CT Based EDS

David Castañón

Boston University

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

D. H. Ye (Marquette)
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?

- 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 S&T
  - Phase 1 (12/17-6/18): Five performers: Boston University, Livermore Laboratory, University of Durham (UK), and two teams from Purdue University
  - Phase 2 (7/18-12/18): Three performers: Boston University, University of Durham, Purdue 1
- 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
  - Many performers changed AATR approach completely from Phase 1 to Phase 2, based on lessons from Phase 1

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 in Phase 1, limited distribution in Phase 2
- Additional objects of interest
  - Several compounds of different sizes, designed by Lawrence Livermore Laboratory
  - Included only in sequestered data

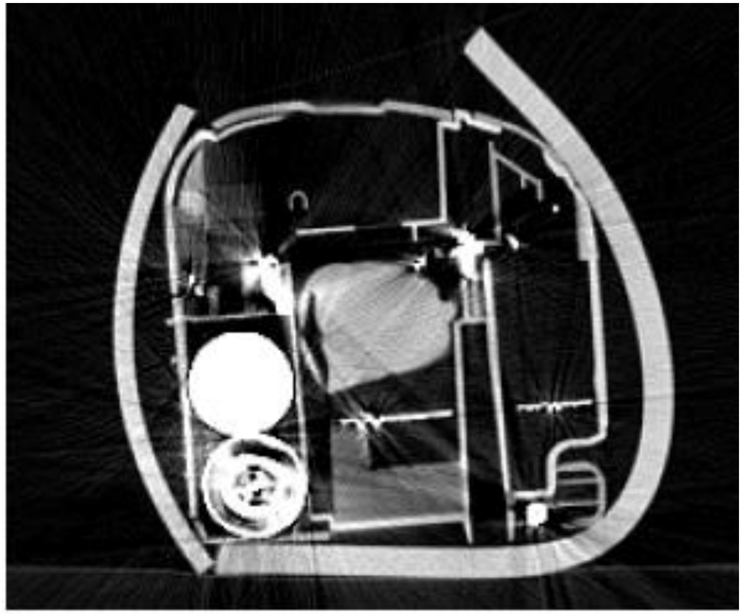

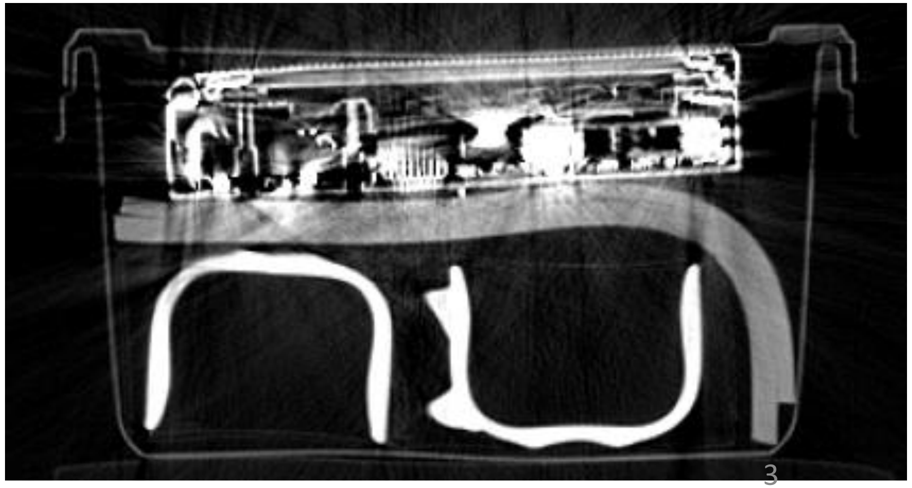

Adaptive ATR Tasks

BOSTON UNIVERSITY

- 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
```

Phase I AATR 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
  - Note: no deep network algorithms given limited size of training data

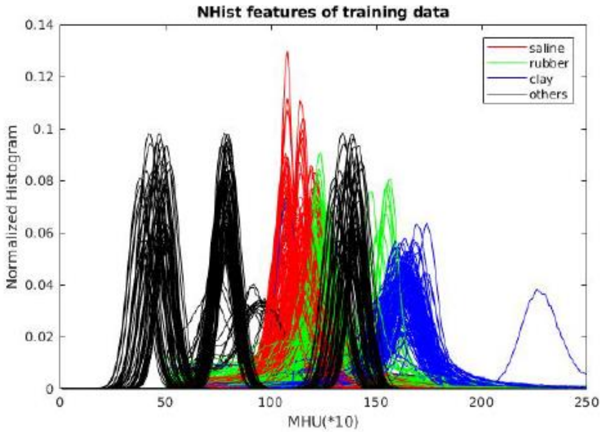

Histogram of intensities/class

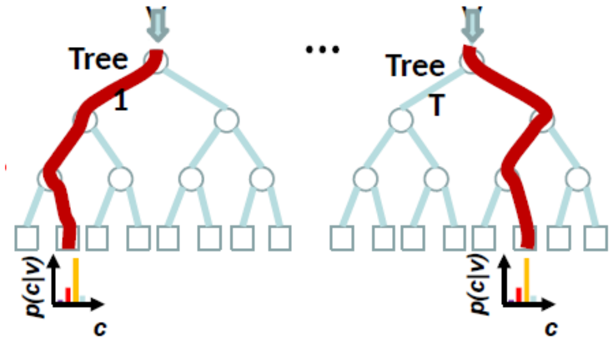

Random Forest Classifier

Evaluation Metrics

- 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

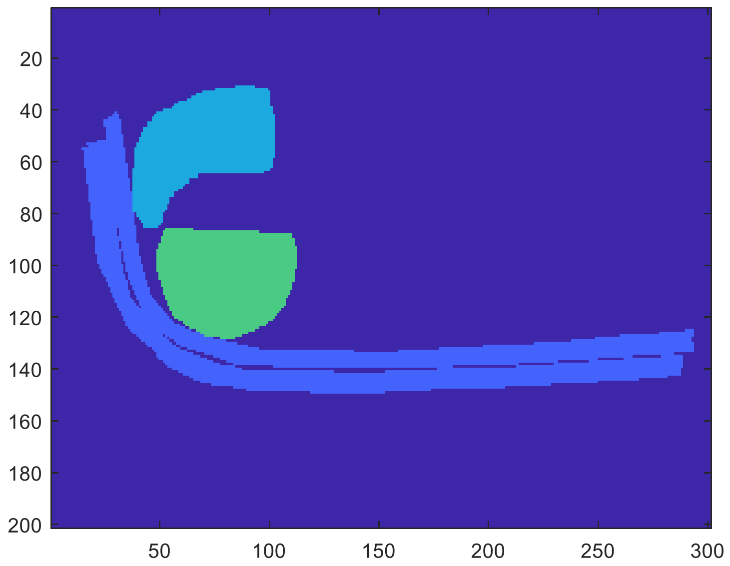

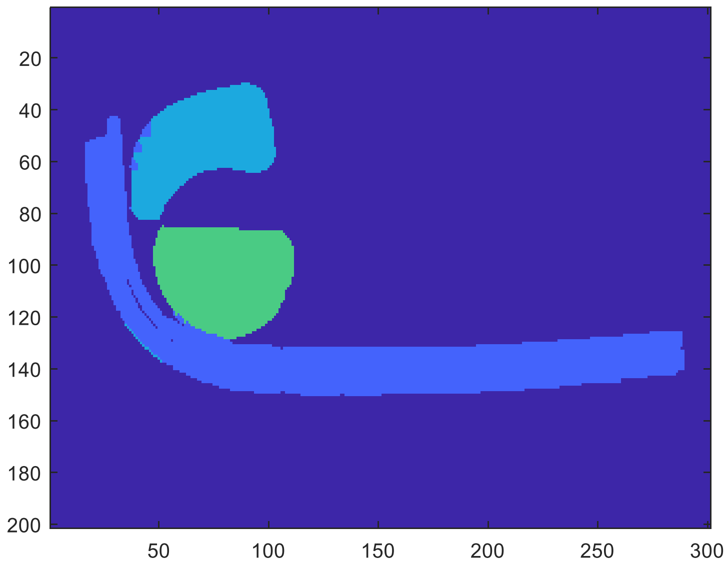

ATR Output

Phase 1: 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)

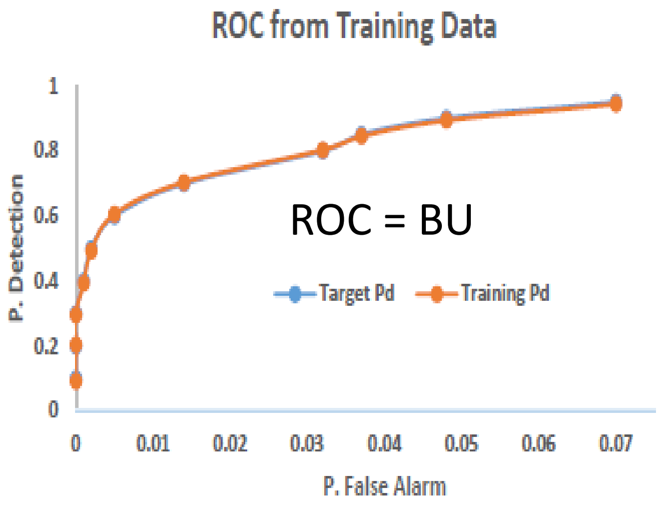

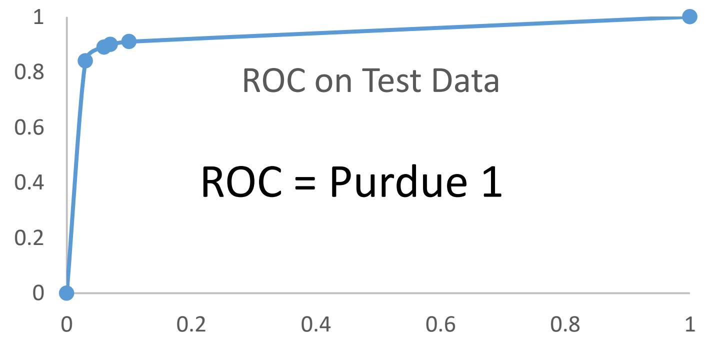

Adaptation to Requirements for Unknown Threats

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 (as in certification tests), but not practical in real scenarios with unknown threats

Phase 1: Performance of Approaches for Unknown Threats

BOSTON UNIVERSITY

- Phase 1: only threat is the unknown threats
  - Four different tests, different unknown materials
- Dichotomy: Teams 1-4 used same ATR features as for known threats; team 5 used parallel classifier
  - Teams 1-4 had to synthesize training features from textual specifications
- Performance of separate classifier (team 5) with reduced feature set as good or better than alternatives.

Illustration of parallel ATR structure (team 5)

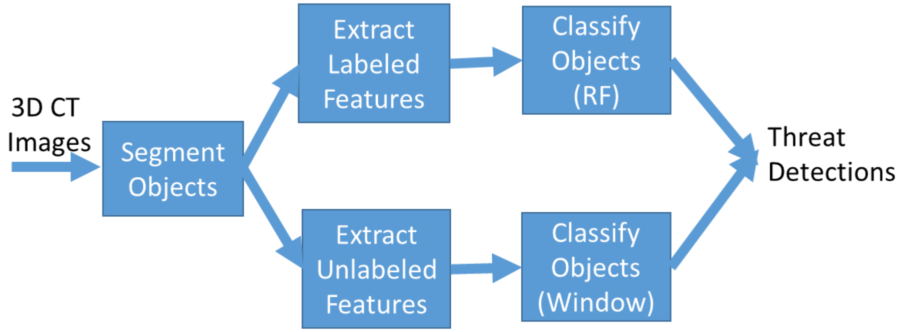

					Durham		Purdue 1		LLNL		Purdue 2		BU	
	Density Range	Minimum	Req.	Req.	Team 1		Team 2		Team 3		Team 4		Team 5	
OOI	(MHU)	Mass (g)	PD (%)	PFA (%)	PD	PFA	PD	PFA	PD	PFA	PD	PFA	PD	PFA
A1	380-525	42	90	10	76	12	83	14	94	11	26	13	89	1
A2	770-810	67	90	10	100	46	100	13	85	4	71	47	100	5
A3	1300-1375	174	90	10	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

Phase 2 Objectives

BOSTON UNIVERSITY

- Test on scenarios with objects of interest that include both new threats as well as threats in the training set
  - Objects of interest include clay (with an extended region of responsibility: +300 MHU in range) as well as other unknown objects
- Evaluate improvement in AATR performance on unknown threats when limited amounts of new training data is provided (from sequestered data)
  - Can develop ATR process on unknown threats
  - Compare performance of AATR with no additional training data vs ATR with additional training data
- Methods: Performers allowed to iterate on algorithm parameters after seeing performance results
  - 4 iterations for AATR algorithm; 3 iterations for ATR algorithm

Phase 2 Results

BOSTON UNIVERSITY

- Desired: P_D 90%, P_{FA} 10%
- Durham abandoned phase 1 approach, treated clay and M5 as materials with single feature: MHU
  - ATR uses more information on segmentation
- Marquette generated synthetic features (histograms) for m5
  - ATR used histograms from limited training data for M5 and Clay, single classifier based on histogram feature
- BU used parallel classifiers
  - Features for clay extended with simulated samples for expanded RoR
  - ATR used more features than AATR for M5, trained on limited data

BU

ORS	PD/PFA	AATR PD/PFA (%) – Iterations				ATR PD/PFA (%) - Iterations		
		1	2	3	4	1	2	3
Clay	PD, clay	90	90			90	90	
	PFA	3	3			3	3	
m5	PD, m5	87	87			60	87	
	PFA	20	5			2	3	
Clay + m5	PD, clay	90	90			90	90	
	PD, m5	87	87			53	87	
	PD, clay+m5	89	89			75	89	
	PFA	24	8			4	6	

Marquette

ORS	PD/PFA	AATR PD/PFA (%) – Iterations				ATR PD/PFA (%) - Iterations		
		1	2	3	4	1	2	3
Clay	PD, clay	67	48	76	86	86	86	86
	PFA	4	2	4	6	10	8	5
m5	PD, m5	100	100	100	100	93	93	93
	PFA	40	38	35	16	11	15	10
Clay + m5	PD, clay	57	48	71	86	86	86	76
	PD, m5	100	100	100	100	87	93	87
	PD, clay+m5	75	69	83	92	86	89	81
	PFA	46	40	41	22	18	22	10

Durham

ORS	PD/PFA	AATR PD/PFA (%) – Iterations				ATR PD/PFA (%) - Iterations		
		1	2	3	4	1	2	3
Clay	PD, clay	95	95	95	95	95	95	95
	PFA	14	3	5	8	7	7	7
m5	PD, m5	100	100	100	87	80	93	92
	PFA	69	46	37	14	20	15	15
Clay + m5	PD, clay	95	52	86	95	100	90	95
	PD, m5	100	100	100	87	87	93	93
	PD, clay+m5	97	72	92	92	94	92	94
	PFA	85	52	44	23	34	28	26

So what was learned?

BOSTON UNIVERSITY

- ATR algorithms can be extended to modify performance in response to computer-readable detection requirement specifications
  - Does require access to training data to tune/evaluate adaptation
  - ALERT developed methods to test AATRs to assure that they are adaptable
- The PD of an AATR is equivalent to an ATR, albeit with a higher probability of false alarm (PFA).
  - When a few training images are supplied to the AATR the PFA of an AATR can be improved to approach the PFA of an ATR while maintaining good PD.
- Difficulty for AATR: How to validate/verify detection performance against new threats
  - 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

Status and Next Steps

BOSTON UNIVERSITY

- TSA and vendors are recognizing importance of AATR development
  - Rapid response to emerging threats
  - Risk-based screening
  - …
- Next steps
  - Develop AATR variations of existing vendor algorithms for EDS
  - Develop process for certification of AATR systems
  - Develop extensions of AATR concepts applicable to checkpoint CT, AT2 and AIT
  - …