AATR Technical Review: March 22, 2018

BOSTON UNIVERSITY

- Institution: Boston University
- Team Members
  - David Castañón, Professor Electrical and Computer Engineering, Systems Engineering
  - Trent Montgomery, Graduate Student, Systems Engineering
- Contact Information:
  - dac@bu.edu, montgomt@bu.edu

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

Executive Summary

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
- Approach:
  - Modify ATR performance by retraining with reweighing
  - Cross-validate predicted performance with training data before deploying
  - For objects with no training data, use hierarchical classifier approach trained based on false alarm performance

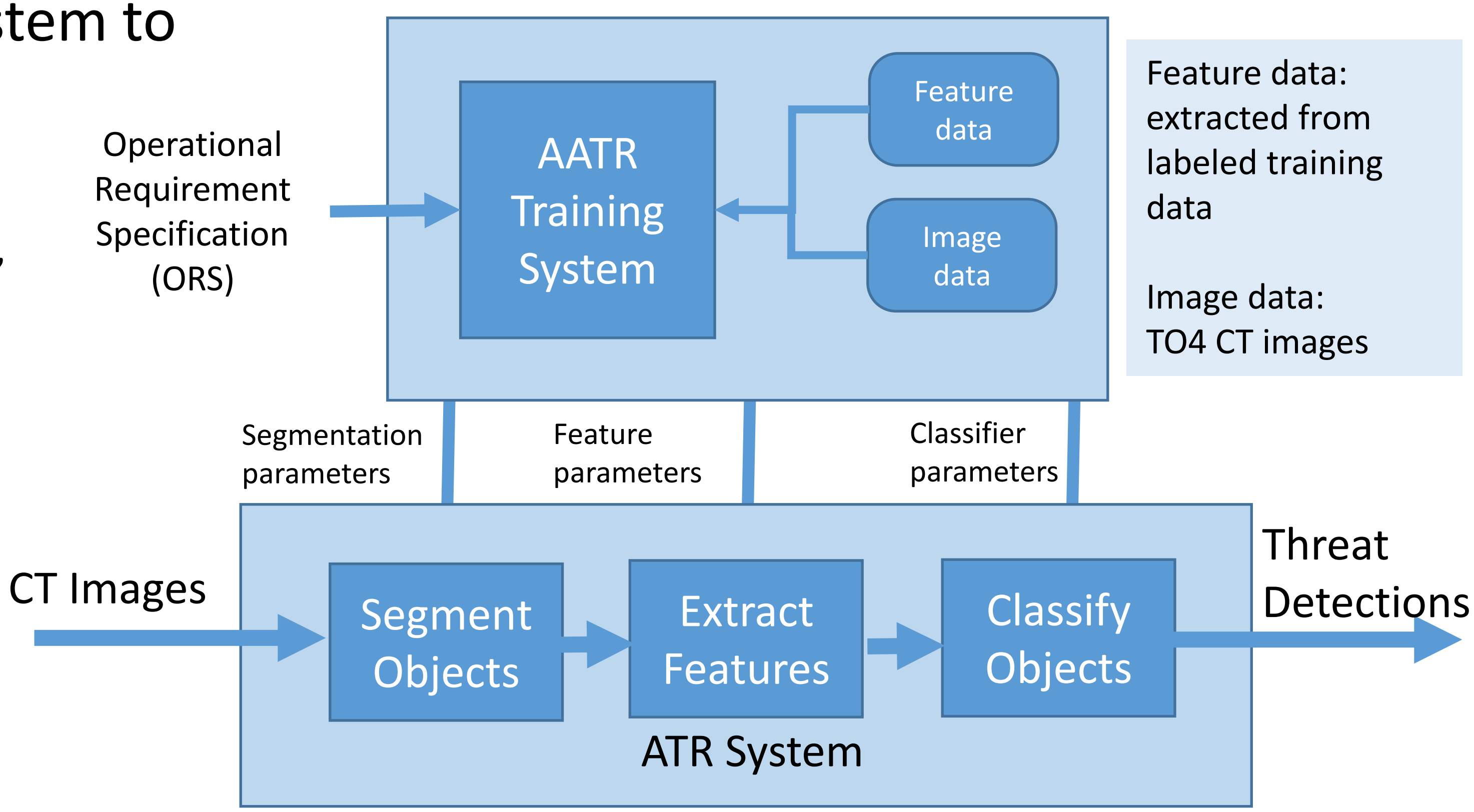

Technical Description: Baseline ATR

BOSTON UNIVERSITY

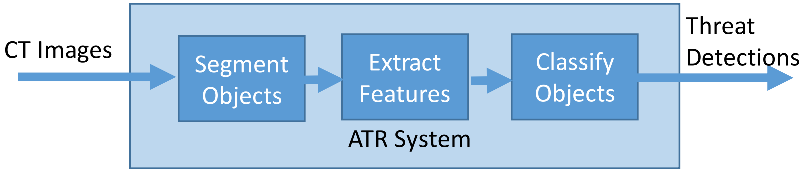

- Standard processing sequence: Segment Objects, Extract features, Classify Objects

- Segmentation
  - Restrict images to densities in window of interest
  - Precondition images with bilateral filtering to enhance contrast, reduce effect of metal artifacts
    - No other metal artifact correction techniques used
  - Connected component labeling for initial segmentation – result in merged and split objects
  - Merged objects initially split along histograms
  - Secondary object splitting technique based on histogram of widths to detect merging of sheet/bulk objects, followed by sheet segmentation
    - Extensions of skeletonization and distance transforms

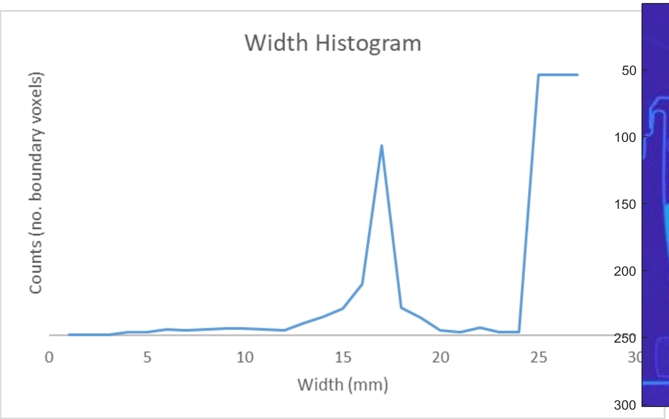

CT slice

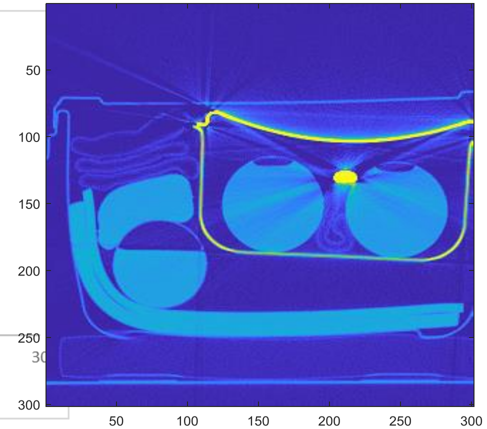

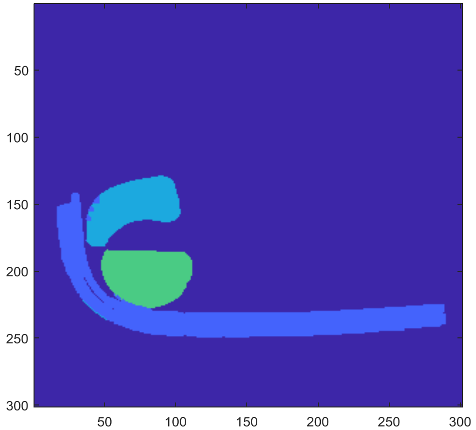

Spit of merged object into sheet, two bulks

Baseline ATR - 2

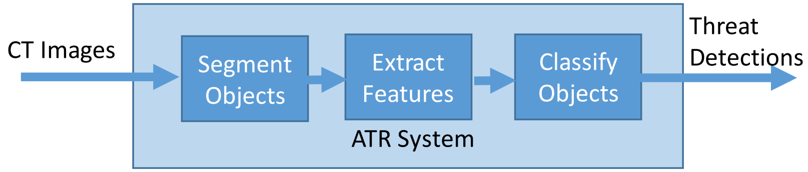

BOSTON UNIVERSITY

- Feature extraction
  - Ground-truth labeled data (TO-4): Segmented volumes with object labels for potential threats
  - For each volume, focus on interior region to extract features
  - Main features: average density, overall density variance, local texture (average variance of density over local textures)
  - Secondary features: volume, mass, thickness, shape (discrete: sheet vs bulk)
  - Note: Secondary features known for many training data objects, must be estimated from data for test objects

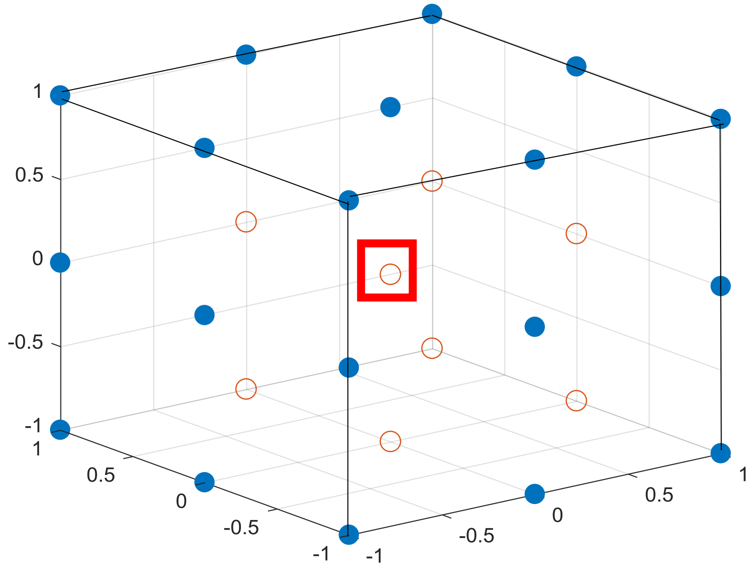

Texture measure:
Standard deviation of density values, averaged across patches in object

Baseline ATR - 3

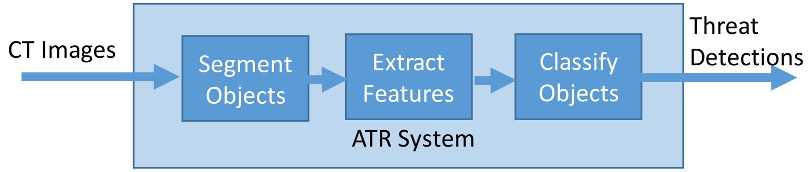

- Classification
  - Two-stage classification: Identify potential threats with primary features, then prune with selected secondary features (mass, thickness)
  - Training data selected based on ORS specification to have balanced threats, confusers in desired region of interest
- Primary Classifier family: Random Forests
  - Ensemble classifier with multiple decision trees (60)
  - Each tree trained with random subset of data, random selection of feature orders
  - Best performer (but not by much) in recent JMLR article (Fernández '14)
  - Believed as more robust to generalization
- Secondary classifier family
  - Minimum size thresholds (mass, thickness)

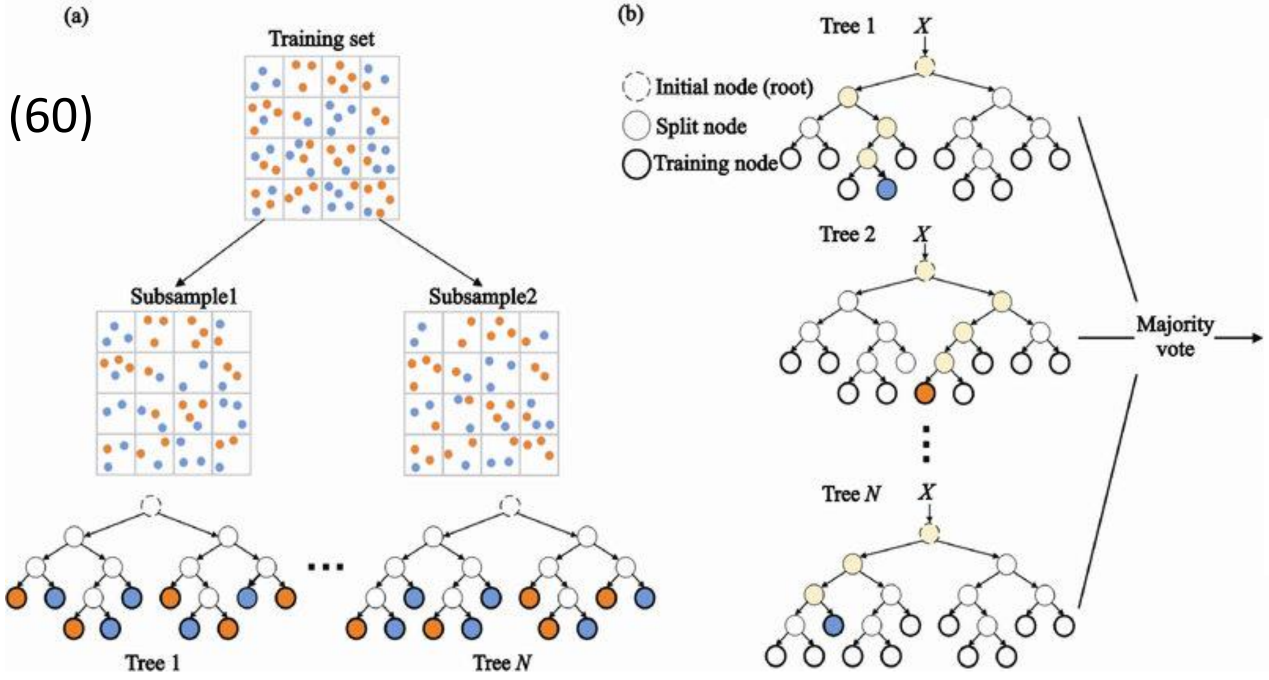

From Machado, Mendoza, Corbellini '15

Adaptive ATR

- Philosophy of approach:
  - Develop general adaptation approaches that are applicable across ATR systems (e.g. SVMs, decision trees, random forests, AdaBoost)
  - Exploit theory for adaptation of parametric classifiers in decision theory

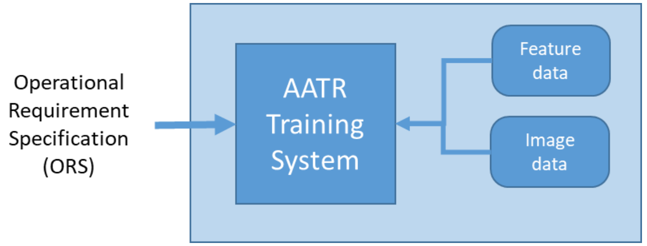

- Adaptation with respect to desired PD/PFA tradeoff (AM1)
  - Parametric approach: modify costs of PD vs PFA errors
  - Our approach: Zadrozny, Langford, Abe '03 – Transparent Box approach: reweight data
  - Select weights using cross validation on training data with bagging: train on 80%, test on 20%

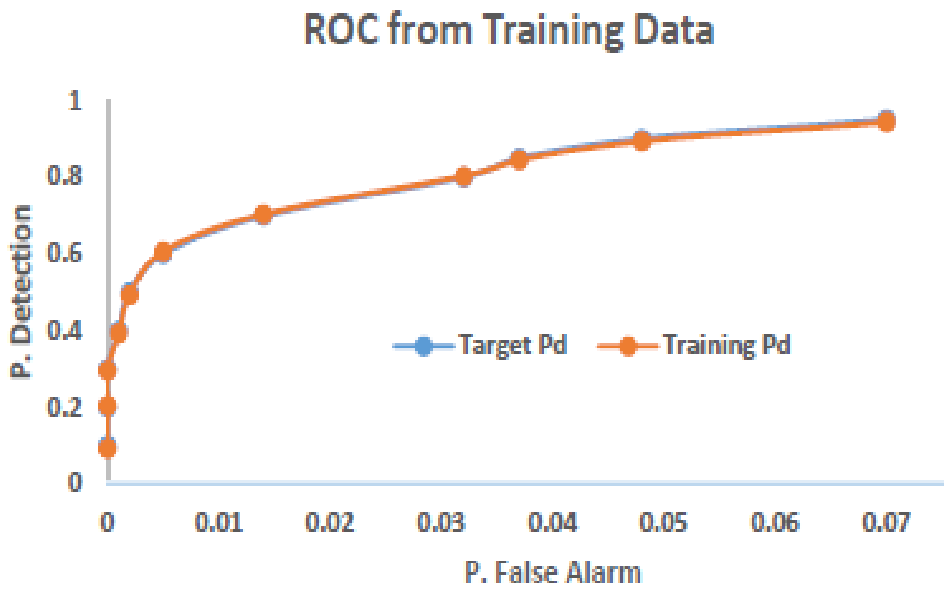

- Results in predicted ROC curve

Adaptive ATR - 2

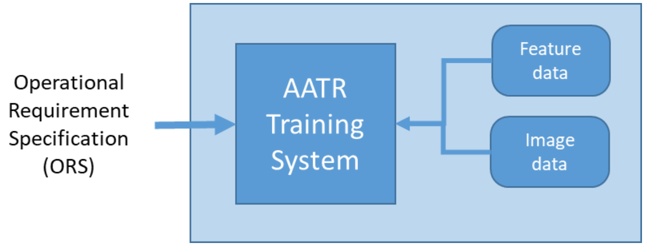

- Adaptation with respect to change of threat material (AM2)
  - Retraining classifier trees
- Adaptation with differential PD/PFA for different materials (AM3)
  - Limited parametric theory: Differential weights on errors for different materials
  - Select differential weights using cross validation on training data
- Adaptation with respect to minimum mass (AM4)
  - Secondary classification: potential threats identified reduced using mass-based threshold
  - Threshold tuned to preserve PD on training data, compensate for segmentation errors
- Adaptation with respect to minimum thickness (AM5)
  - Secondary classification: potential threats reduced using minimum thickness
  - Threshold tuned to preserve PD on training data, allow for errors in thickness estimates

Adaptive ATR - 3

- Adaptation with respect to new threats with no training data (TO7-AM2)
  - Main issue: no training data (labeled or otherwise)
  - Limits ability to extract same features, use similar classifiers
  - Only information provided: density range, minimum mass and thickness
- Approach: Hierarchical classifier
  - Parallel classifiers: baseline as previous plus added classifier using only features as in ORS
  - New features: average density, mass, thickness

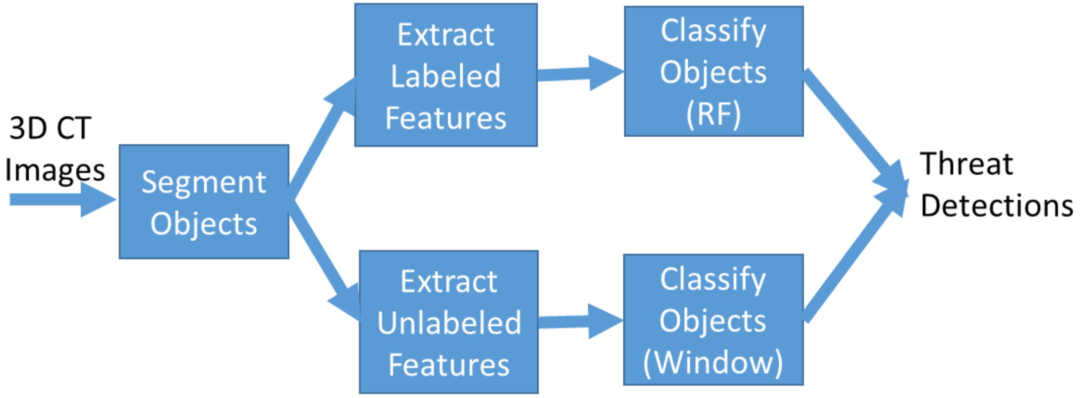

- Cannot use cross-validation to estimate PD for added classifier: hard to predict performance
  - No training samples of threats
  - Select parameters based on training data PFA
  - Alternative: simulate samples (ideally with full tomography model), not used

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

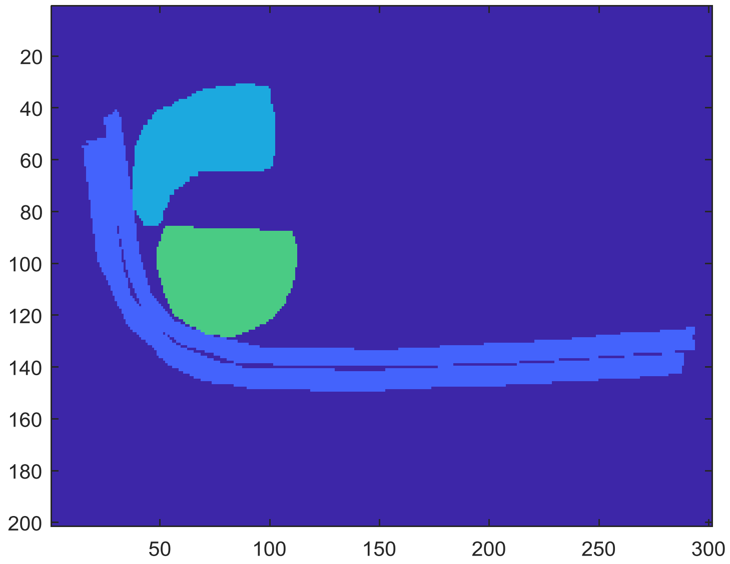

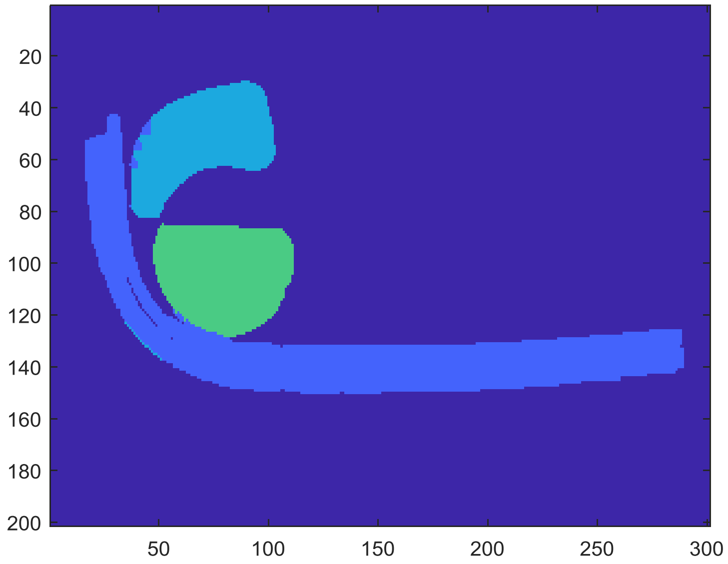

ATR Output

Results & Comments

- Some segmentation leakage limits upside in PD
  - Particularly with limited metal artifact correction
  - Best seen in AM4 where large mass objects are split
- Random forest shows significant variability in performance
  - Insufficient training data for technique?
- Hard to get second PD down in AM 3 with reweighting approach

Performer Training / TO4 Data

AM 1: AROC

OOI	Required PD	Required PFA	AATR PD	AATR PFA
S	0.7	0.02	0.74	0.06
S	0.8	0.05	0.76	0.08
S	0.85	0.08	0.83	0.10
S	0.9	0.1	0.89	0.12
S	0.95	0.2	0.9	0.13
AROC		0.91		

AM 2: PD/PFA for Varying OOIs

OOI	Required PD [%]	Required PFA [%]	AATR PD [%]	AATR PFA [%]
C,S,R	90	10	90	20
C	90	10	93	20
S	90	10	92	20
R	90	10	86	20

AM 3: Varing PD Weight

OOI	Required PD [%]	Required PFA [%]	AATR PD [%]	AATR PFA [%]
C,S	C:90, S:90	10	C: 85, S:89	15
C,S	C:20, S:90	10	C: 34, S:89	14
C,S	C:90, S:20	10	C: 85, S:43	11

AM 4: PD/PFA for Varying Mass

OOI	Min Mass [g]	Required PD [%]	Required PFA [%]	AATR PD [%]	AATR PFA [%]	Incremental Mass Rnge [g]	AATR Incremental PD [%]
S	400	90	10	82	10	N/A	N/A
S	300	90	10	87	11	300 - 400	90
S	100	90	10	91	12	100 - 300	92

AM 5: PD/PFA for Varying Thicness

OOI	Min Thickness [mm]	Required PD [%]	Required PFA [%]	AATR PD [%]	AATR PFA [%]	Incremental Thickness Rnge [mm]	AATR Incremental PD [%]
R	10	90	10	80	9	N/A	N/A
R	6.5	90	10	74	10	6.5 – 10	71
R	0	90	10	70	10	0 – 6.5	62

ALERT Testing / TO7 Data

AM 2: PD/PFA for Varying OOIs

OOI(s)	Density Rnge [MHU]	Minimum Mass (g)	Required PD [%]	Required PFA [%]	AATR PD [%]	AATR PFA [%]
m1	380-525	42	90	10	89	1
m2	770-810	67	90	10	100	5
m3	1300-1375	174	90	10	92	11
m4	1350-1430	183	90	10	100	0

Discussion

- AATR approach with reweighting retraining adapts well to range of ORS specified
  - Performance can be improved with enhanced metal artifact reduction, segmentation
  - Extensible across classifier families
- Hierarchical classifier for new threats is practical, performs well
  - May be "interim" solution for fieldable system without perturbing baseline ATR
  - Need to solve performance verification gap – Perhaps through use of enhanced simulation
- Areas for improvement
  - Enhanced metal artifact correction/compensation – e.g. detect nearby metal as a feature and adapt classifier
  - Additional feature selection (without overtraining)
  - Classifier families with less randomness in behavior – using reweighting approach
  - Extension of reweighting approach to integrate both PD, PFA targets
  - Exploration of performance when many classes of threats are present – how to control PFA