Proposer's Day: BIG-R BAA Topic 8 (SPECTRE)

Jacob McCord, Triet Le PhD, Chris Algire, Macarena Ortiz PhD

Agenda

11:05 EDT	NGA Research Overview – Dr. John Main
11:25 EDT	SPECTRE Program Overview
12:30 EDT	SPECTRE Timeline, Expectations, and Logistics
1:30 EDT	Lightning Talks
2:20 EDT	Q&A Session
3:00 EDT	Closing Remarks

Please ask all questions via chat

Keynote – NGA Research Overview
Dr. John Main (Acting Deputy Director, NGA Research)

SPECTRE Program Overview

Computer Vision Today

Time consuming, costly, risky

Computer Vision Tomorrow

Build on Weak Supervision Innovation

Computer vision traditionally uses **strong supervision**, not **weak supervision**

STRONG *(TIME & COST INTENSIVE)* → **WEAK** *(TIME & COST EFFICIENT)*

12 months per object → *2 months per object*

Less Abstract

Harder to Create

More Abstract

Easier to Create

Pixel-level Labels | Bounding Box Labels | Point Labels | Count Labels | Caption Labels

Images drawn from NGA's xView dataset

Bridge the Gap between Weak Supervision and GEOINT

Goal: Combine computer vision and natural language processing to achieve **weakly supervised object detection** in overhead imagery

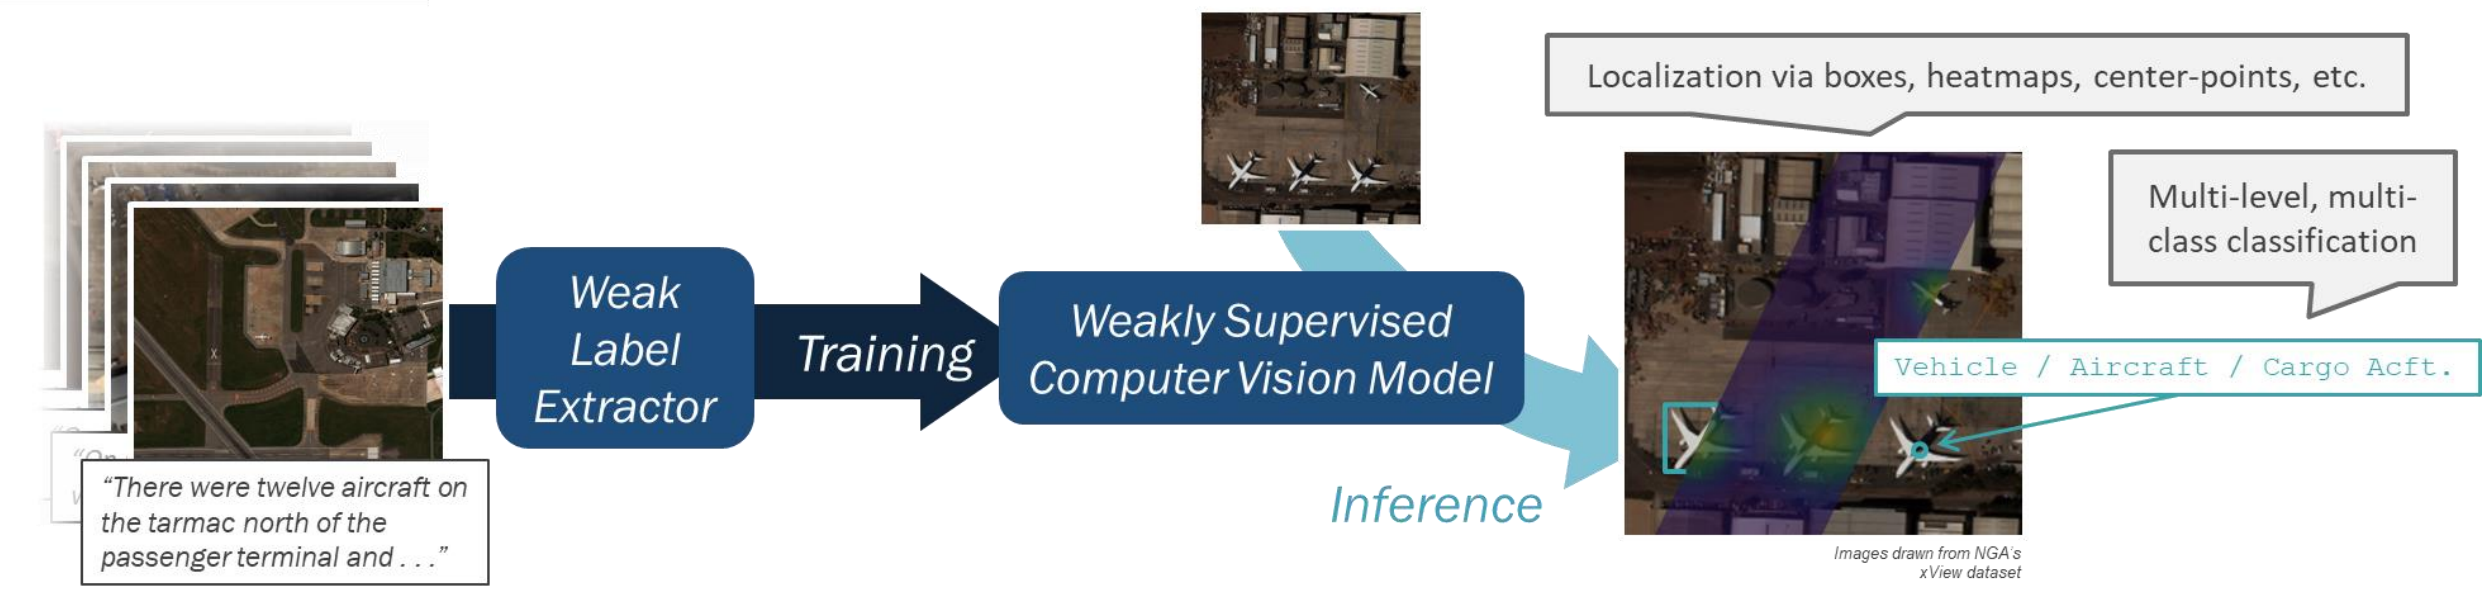

Current State-of-the-Art Weak Supervision

Manual Labeling

(Includes active learning, semi-supervised and supervised learning at the image level)

Images drawn from NGA's xView dataset

Reliant on subject matter experts

Current State-of-the-Art Weak Supervision

Heuristic Labeling Functions and Rule-based Knowledge Bases

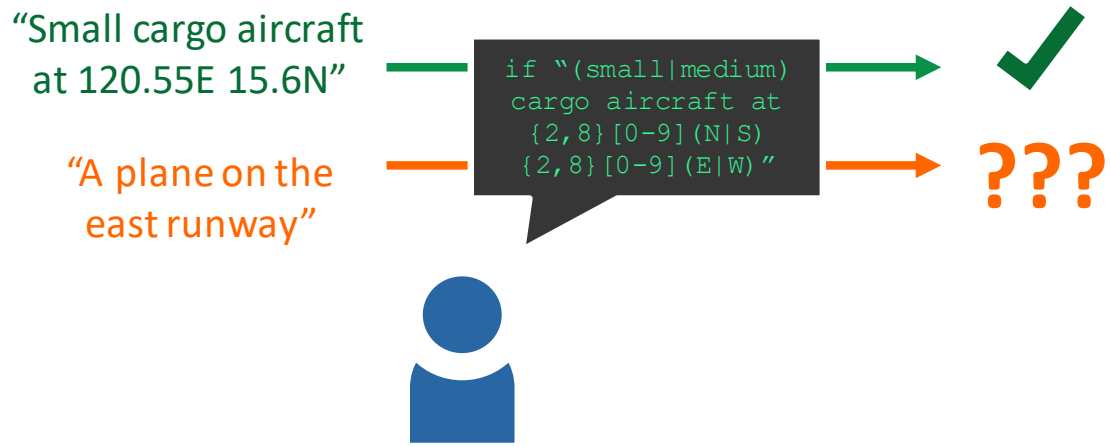

Reliant on subject matter experts and
requires noiseless and comprehensive text

Current State-of-the-Art Weak Supervision

Existing Embedding/Attention/Saliency Image-Level Methods

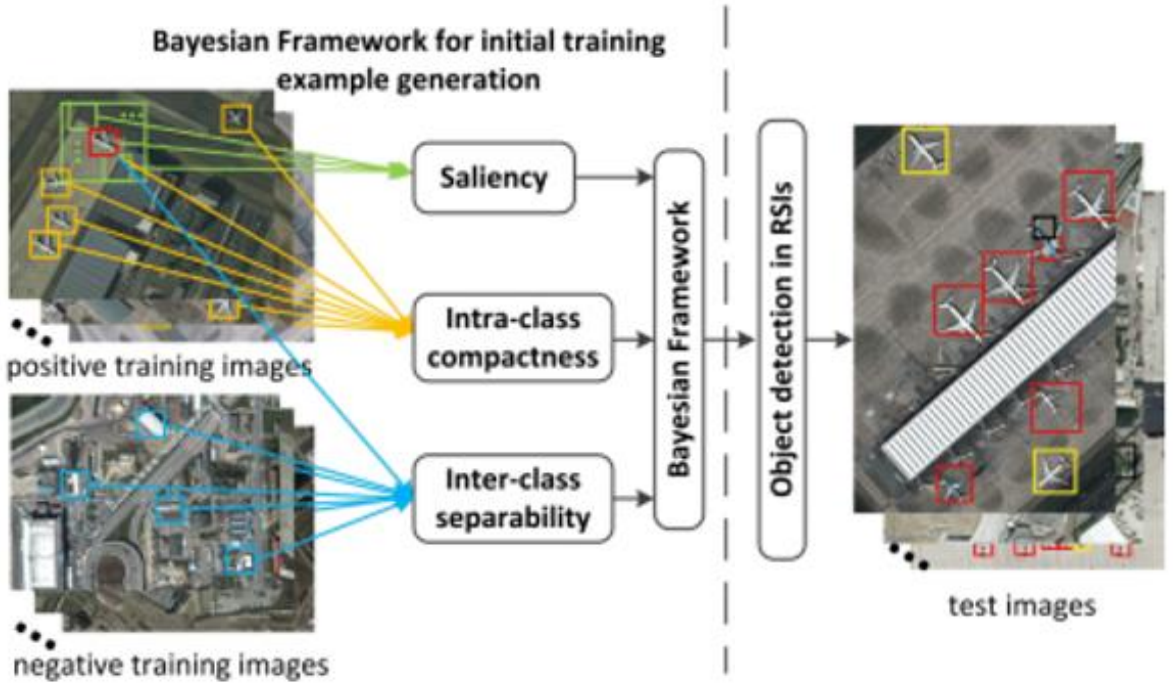

J. Han, D. Zhang, G. Cheng, L. Gu and J. Ren, "Object Detection in Optical Remote Sensing Images Based on Weakly Supervised Learning and High-level Feature Learning," 2015.

Binary or n-way multi-class classification
***ignores semantic relationships** of classes*

Weak Supervision with GEOINT: Needs Innovation Beyond State-of-the-Art

Two medium (possibly Mi-8 or Mi-14) transport helicopters were parked at 39.1576N, 127.4903E and 39.15 …

Several small vehicles were parked on the tarmac near the MiG-29 and Il-28 aircraft.

© 38 North - https://www.38north.org/2019/11/wonsan111419/

Training

Weakly Supervised Computer Vision Model

GEOINT describes images in unique ways: this presents opportunities and challenges that require novel weak supervision methods

Innovation Required in Several Technical Elements

Foundations of Natural Language Processing

- Extracting weak supervision information from text and fusing it with imagery

Foundations of Computer Vision

- Detecting and classifying objects in large-scale overhead imagery

Extensions to Remote Sensing

- Adapting weak supervision to sources of noise found in GEOINT

All proposals should address all technical elements

Key Challenges of Weak Supervision with GEOINT

Two **medium (possibly Mi-8 or Mi-14) transport helicopters** were parked at **39.1576N, 127.4903E and 39.15** …

Several **small vehicles** were parked **on the tarmac near** the **MiG-29 and Il-28** aircraft.

© 38 North - https://www.38north.org/2019/11/wonsan111419/

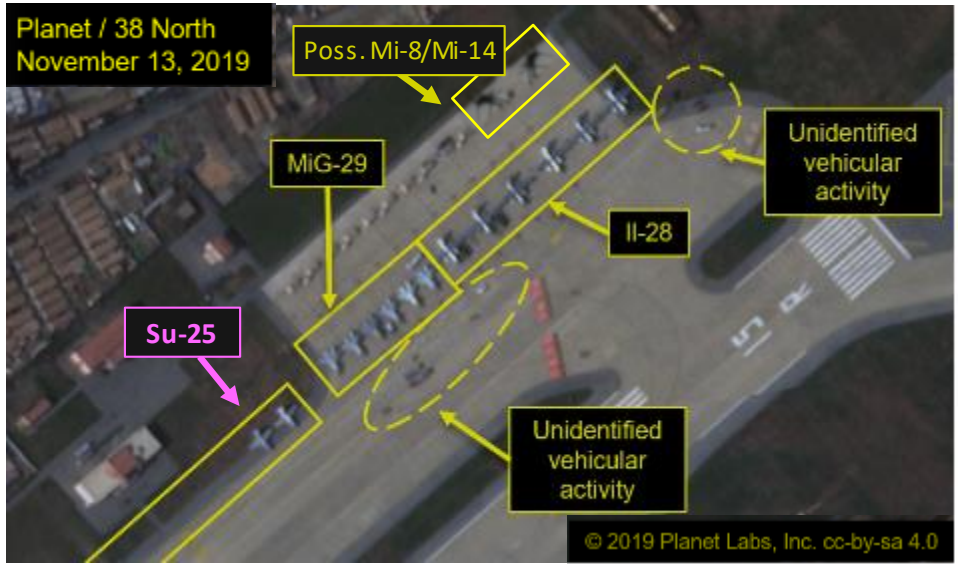

Location - Class - Completeness

Foundations of CV & NLP

Extensions to Remote Sensing

Location – Large-Scale Imagery and Weak Supervision

(Technical Element: Foundations of Computer Vision)

Overhead imagery is large. Each object is often <0.5% of the total pixels.

Image drawn from Wikimedia Commons

ImageNet

Image drawn from NGA's xView dataset

xView 1

Location – Large-Scale Imagery and Weak Supervision

(Technical Element: Foundations of Computer Vision)

Many state-of-the-art algorithms tile imagery, others try to process at the image level. Each has trade-offs.

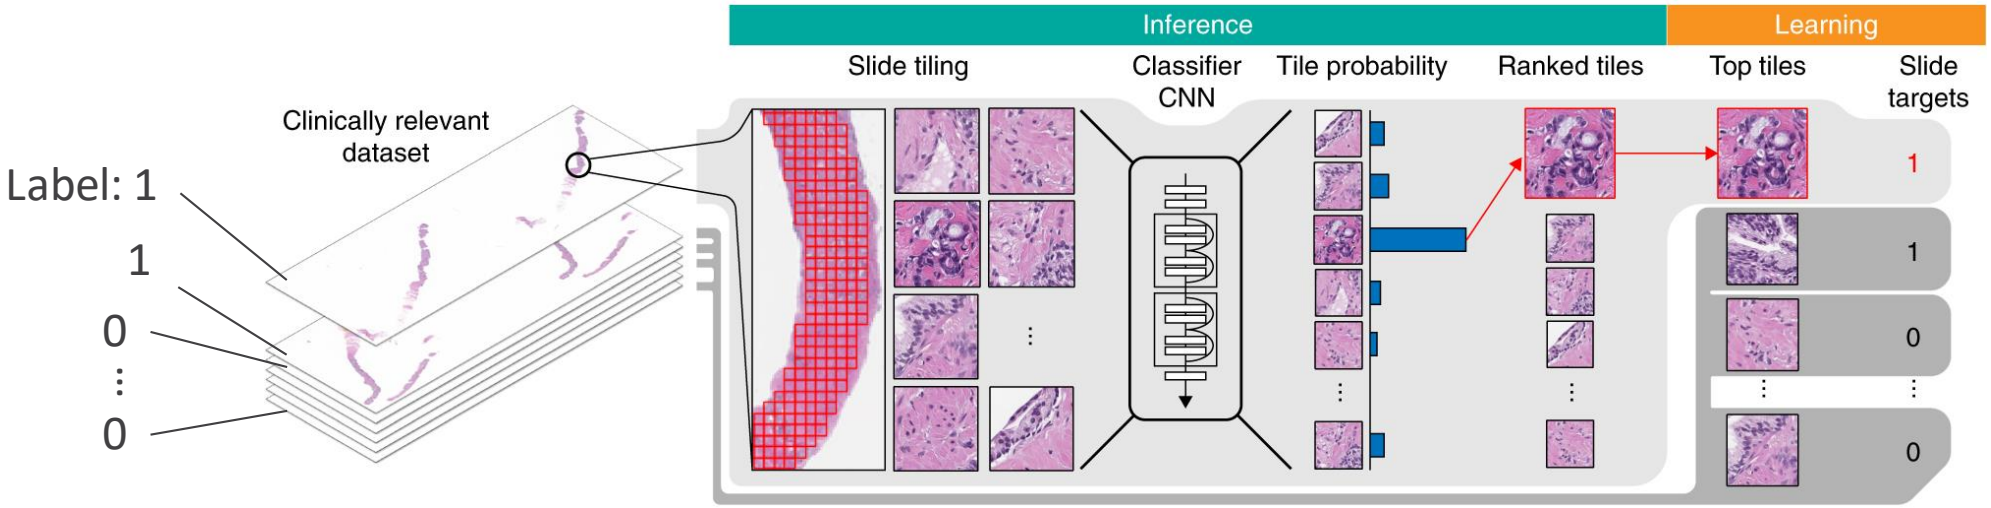

G. Campanella, M. Hanna, L. Geneslaw, A. Miraflor, V. Werneck Krauss Silva, K. Busam, V. Reuter, D. Klimstra, T. Fuchs and E. Brogi, "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images," Nature Medicine, vol. 25, pp. 1301-1309, 2019.

Enhance weak supervision by fusing image and text to guide localization

Location – Fusing Text and Imagery for Weak Supervision of Obj. Det.

(Technical Element: Foundations of Computer Vision, Foundations of Natural Language Processing)

Two **medium (possibly Mi-8 or Mi-14) transport helicopters** were parked at **39.1576N, 127.4903E and 39.15** …

Several **small vehicles** were parked **on the tarmac near** the **MiG-29 and Il-28** aircraft.

© 38 North - https://www.38north.org/2019/11/wonsan111419/

How to represent location relations in image space?

Weak Supervision → *Computer Vision Model*

Challenge: Use ambiguous location information to supervise an object detector

Class – Weak Labels do not have Standardized Classes

(Technical Element: Extensions to Remote Sensing)

Human captions (especially GEOINT) will use umbrella terms, aliases, etc.

Literature exists for hierarchical classification. Dealing with class inconsistency in weak supervision is an open area of research.

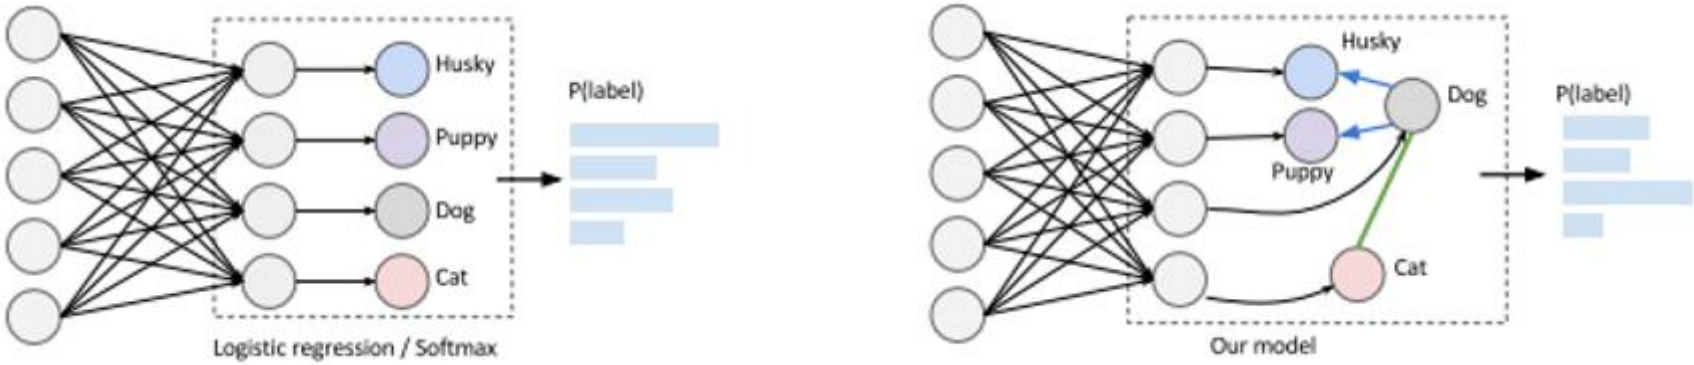

J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven and H. Adam, "Large-scale object classification using label relation graphs," in European conference on computer vision, 2014.

Class – Use Semantic Knowledge to Maximize Training Data

(Technical Element: Extensions to Remote Sensing)

Two **medium (possibly Mi-8 or Mi-14) transport helicopters** were parked at **39.1576N, 127.4903E and 39.15** …

Several **small vehicles** were parked **on the tarmac near** the **MiG-29 and Il-28** aircraft.

© 38 North - https://www.38north.org/2019/11/wonsan111419/

Weak Supervision

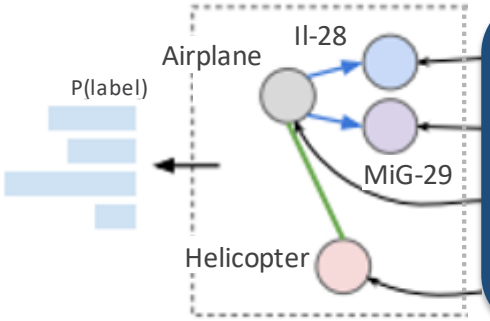

Computer Vision Model

Challenge: Adapt to class noise (given a known object hierarchy)

Completeness – Weak Supervision without Negative Labels

(Technical Element: Extensions to Remote Sensing)

GEOINT is created to answer specific analytic requirements. This is a source of noise when objects are relevant to some requirements but not others (i.e., labeled in some imagery and not others).

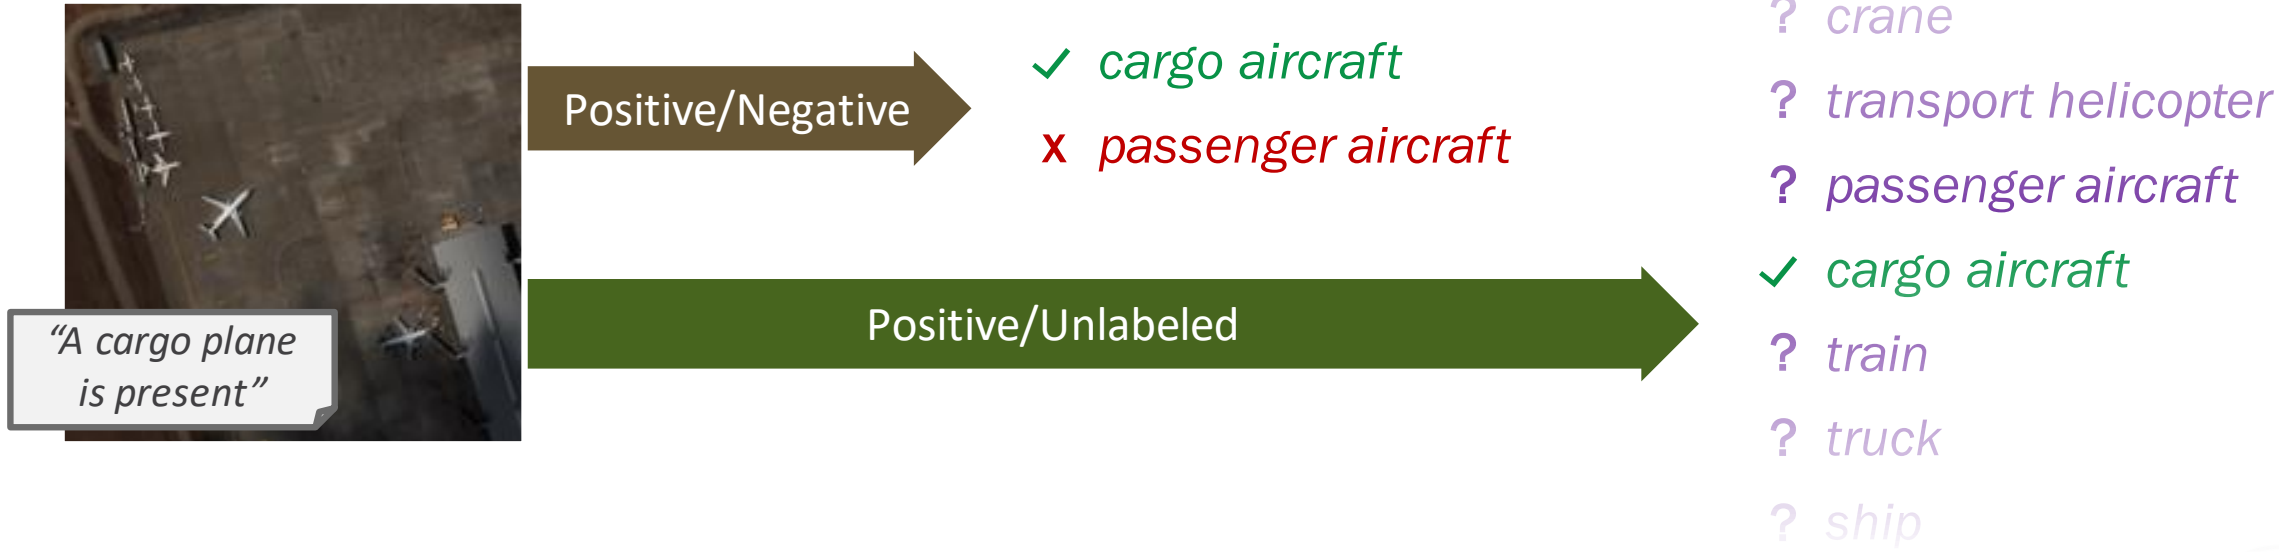

Y. Yang, K. Liang, and L. Carin, "Object Detection as a Positive-Unlabeled Problem." arXiv preprint arXiv:2002.04672.

Completeness – Weak Supervision without Negative Labels

(Technical Element: Extensions to Remote Sensing)

Two **medium (possibly Mi-8 or Mi-14) transport helicopters** were parked at **39.1576N, 127.4903E and 39.15** …

Several **small vehicles** were parked **on the tarmac near** the **MiG-29 and Il-28** aircraft.

© 38 North - https://www.38north.org/2019/11/wonsan111419/

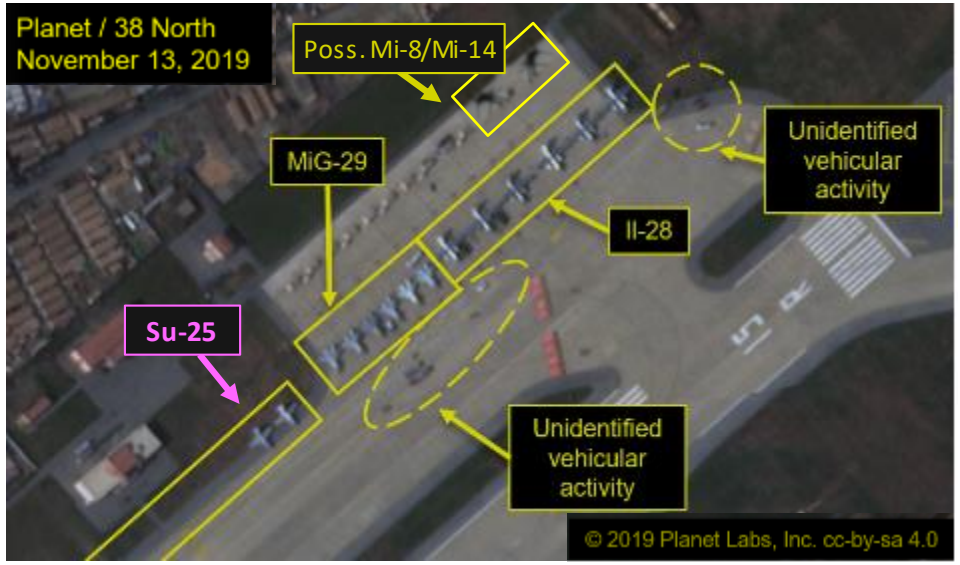

How does this impact the other key challenges?

Challenge: Weak supervision without hard negative labels

These Challenges are Not Exhaustive

Weakly supervised object detection with GEOINT is about designing computer vision models to learn from noisy, incomplete information from unstructured text

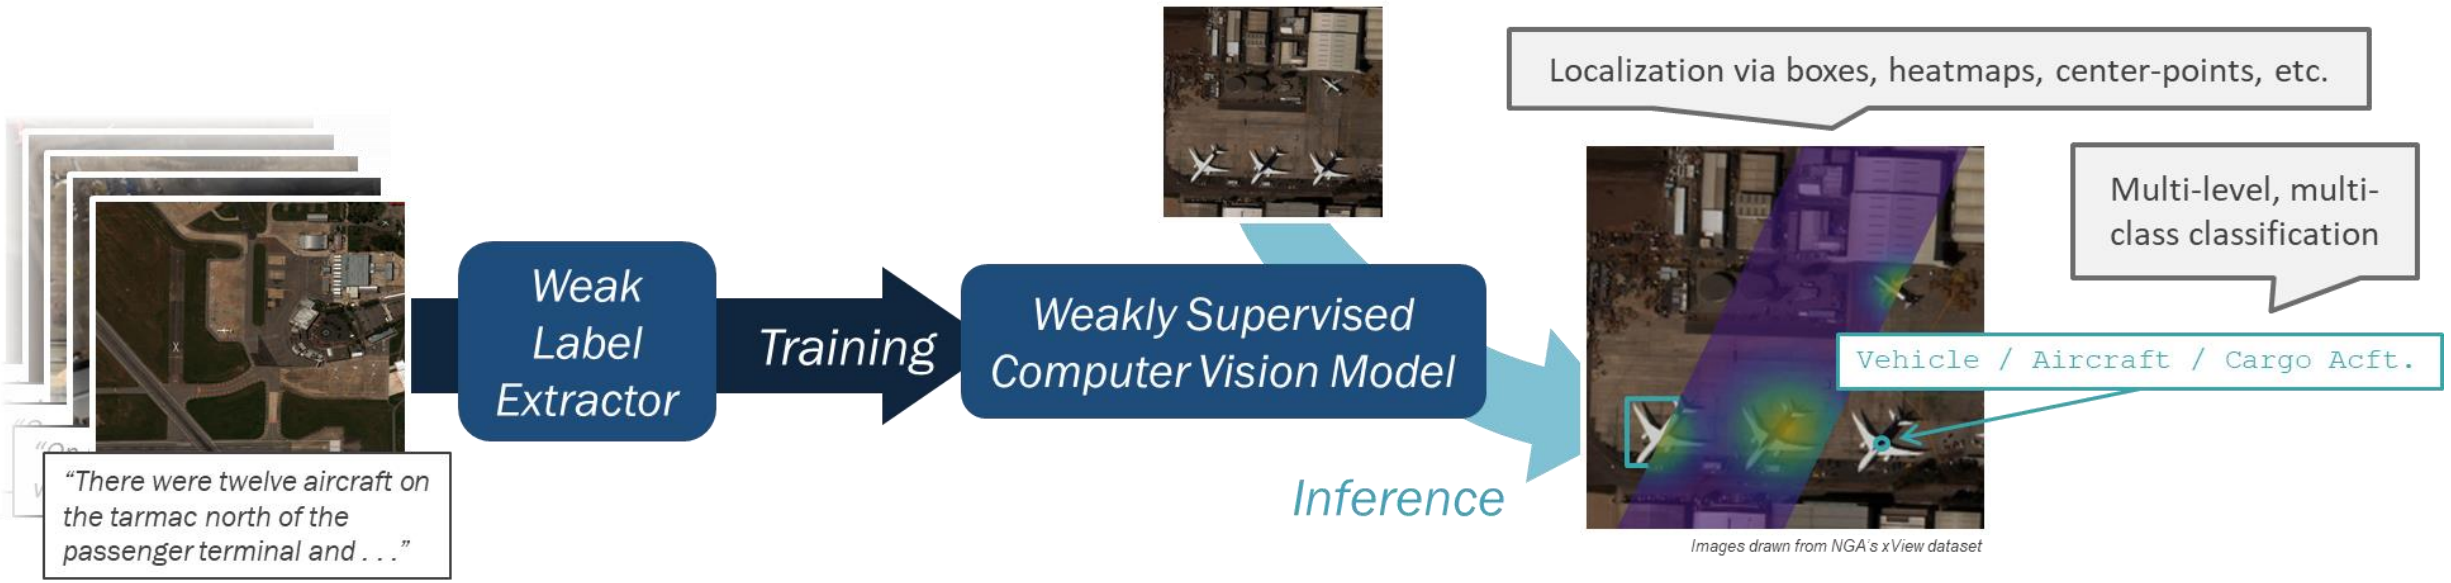

What other innovative solutions exist for fusing GEOINT data through weak supervision?

Break – return at 12:30 P.M. EDT

Program Timeline, Expectations, and Logistics

Overview

Base Period (12 months)

- Develop a Weakly Supervised Object Detection prototype. Address challenges related to the fusion of image and text data (Foundations of CV and NLP).

Option Period (12 months)

- Address challenges related to class and completeness (Extensions to Remote Sensing).

Transition Period (Optional, 6 months)

- Advance technological readiness level of prototype.

All work will be fully UNCLASSIFIED in all periods

The Government retains the right to decide whether or not to exercise the option and transition periods.

Pre-Award Timeline

Topic Release	Questions Due	Q&A Posting	Abstracts Due	Abstract Feedback	Proposals Due
31 May	22 June	7 July	21 July	11 August	6 September

Post-Award Timeline

Base & Option Period

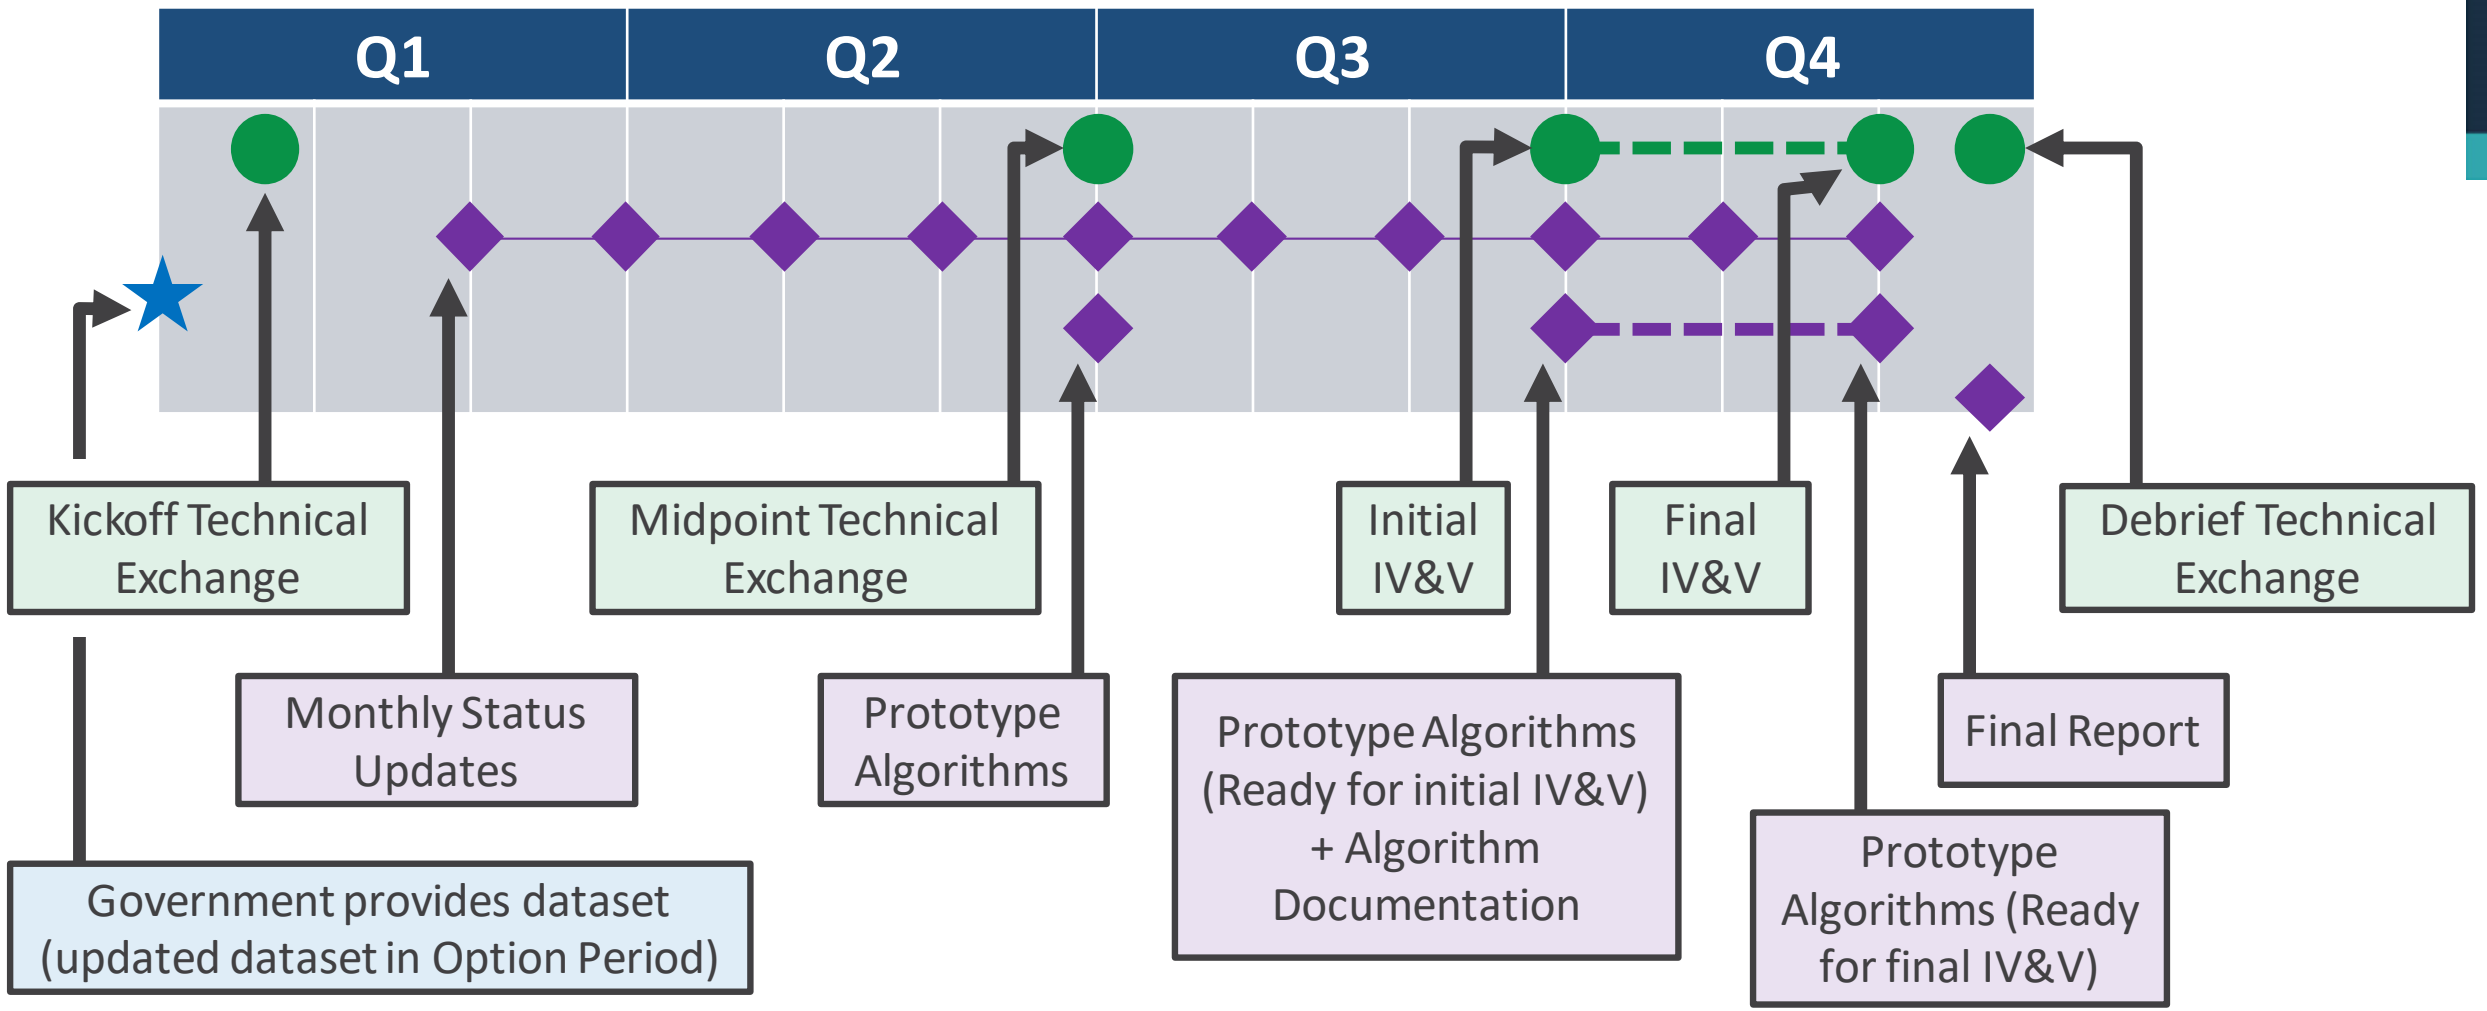

Post-Award Timeline

Transition Period

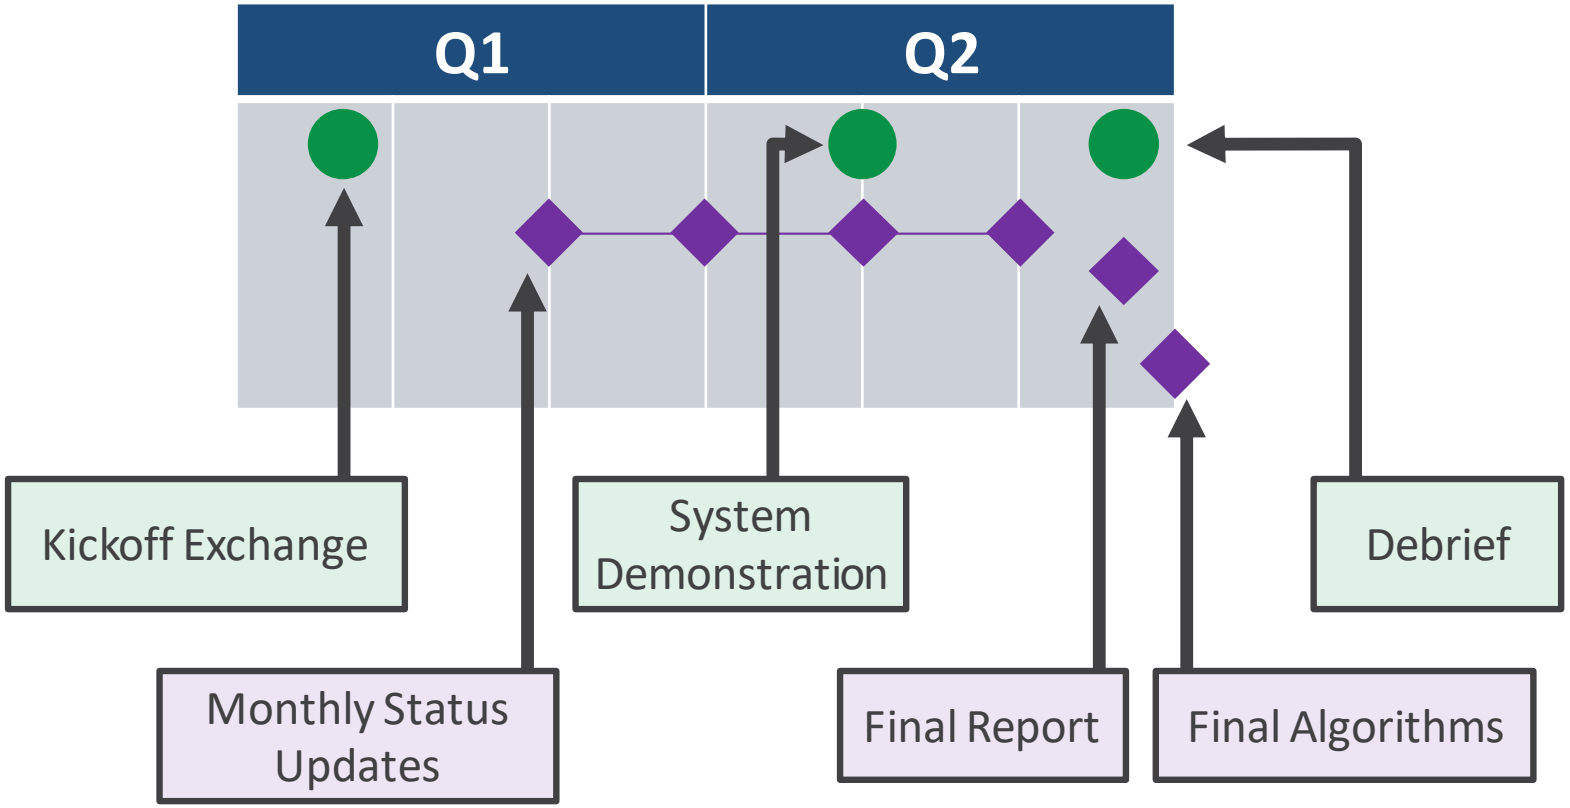

Deliverables

Kickoff Materials	Briefing that details the experimental approach, schedule and milestones.
Monthly Status Reports	Summary of work accomplished, any challenges or issues that may impact cost/schedule/performance, and intended actions for the next reporting period.
Validation and Verification (V&V) Plan	Description of the methodology for evaluating the performance of the proposed solution, and a summary of any required data and code that will be required for evaluation.
Algorithm(s)	Algorithms for the extraction of computer vision labels from text and the use of those labels for object detection. To be delivered in a dockerized container for implementation in a commodity cloud environment.
Algorithm Description Document (ADD)	A detailed ADD and signal model to describe how the detections are created through processing. To be delivered with each algorithm.
Final report	Summary of the work performed and the final results of the investigation.

Government Furnished Information (GFI)

Modification of the xView dataset, with a training split of imagery and corresponding structured information and unstructured text, and

A public validation split with ground-truth object center points and classes

Train

Validation

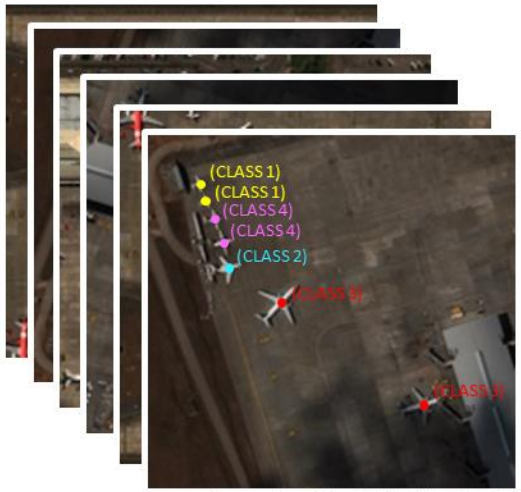

Images drawn from NGA's xView dataset

Government Furnished Information (GFI)

In the Base period, the dataset will reflect technical areas 1 & 2 (Foundations of Natural Language Processing, Foundations of Computer Vision)

- Noise in location and count information

In the Option period, an updated dataset will be provided with additional semantic difficulty, addressing technical area 3 (Extensions to Remote Sensing)

- Noise in class information
- Missing some negative labels
- Taxonomy will be provided

Additional datasets may be proposed (but will only be disseminated with proposer's approval)

Independent Verification and Validation

Initial IV&V test will take place 9 months post-award: allows performers to check progress

Final IV&V test 11 months post-award: evaluation against government expectations

Metrics	Base Period	Option Period
Object Detection	Mean Average Precision (mAP)	
	Threshold: 0.25 Objective: 0.45	Threshold: 0.35 Objective: 0.60
	Recall	
	Threshold: 0.70 Objective: 0.80	Threshold: 0.75 Objective: 0.90
Object Classification	F_1-score	Semantic Distance
	Threshold: 0.70 Objective: 0.80	TBA

Proposals may include additional metrics if proposers wish to demonstrate other areas of performance.

Meeting/not meeting metric thresholds is not a guarantee of selection/non-selection for option periods

Research Transition

Transition will focus on maturing the technological readiness level of algorithms developed in the option period.

This includes adapting algorithms to interface specifications provided by the Government and demonstrating performance on Government-provided datasets.

NGA anticipates that achieving the goals of the program will necessitate a minimum of Government Purpose Rights (as defined in DFARS 227.7103) in all deliverables.

Lunch

Return at 1:25 P.M. EDT for Lightning Talks

Lightning Talks

Break

Return at 2:20 P.M. EDT for Q&A Session

Question Session

Please ask all questions via chat

Questions can also be emailed to **BIGRBAA_T8_Proposers_Day@nga.mil** after the event

Closing Remarks

Topic Release	Questions Due	Q&A Posting	Abstracts Due	Abstract Feedback	Proposals Due
31 May	22 June	7 July	21 July	11 August	6 September

Questions can be emailed to
BIGRBAA_T8_Proposers_Day@nga.mil