DC7119 – MAKING DEEP LEARNING SCALE: DEFENSE APPLICATIONS

WILL RORRER, PROGRAM MANAGER

RON KNEUSEL, PHD, SENIOR DATA SCIENTIST

NVIDIA GTC DC 01 – 02 Nov 2017

HARRIS® TECHNOLOGY TO CONNECT, INFORM AND PROTECT™

Agenda

Harris Corporation Introduction

A Call to Action: The Urgency Behind the DoD's Adoption of AI

Deep Learning Research at Harris

Harris's Work to Scale Deep Learning for Defense

Q&A

Harris Corporation introduction – segment overviews

Communication Systems

Tactical and airborne radios, night vision technology, and defense and public safety networks

Electronic Systems

Electronic warfare, avionics, robotics, advanced communications and maritime systems for the defense industry, as well as air traffic management solutions for the civil aviation industry

Space and Intelligence Systems

Complete solutions encompassing advanced sensors and payloads, processing systems, and analytics for global situational awareness, space superiority missions, and Earth insights

A call to action: the urgency behind the DoD's adoption of AI

"We're going to find ourselves in the not too distant future swimming in sensors and drowning in data"

Lt. Gen. David A Deptula,
USAF Dep Chief of Staff for ISR

2010

"The skies will 'darken' with the hundreds of small satellites to be launched by U.S. companies and as procedures are developed to allow safe operation of unmanned aerial vehicles in civil airspace,"

Robert Cardillo,
Director – NGA

2015

"So just how big is this rising tide? If we were to attempt to manually exploit the commercial satellite imagery we expect to have over the next 20 years, we would need eight million imagery analysts. ***Even now, every day in just one combat theater with a single sensor, we collect the data equivalent of three NFL seasons – every game. In high definition!***

Imagine a coach trying to understand the strategy of his opponents by watching every play made by every team in every game for three seasons – all in one single day. Because three more seasons will be coming tomorrow. That's what we ask our analysts to do – when we don't augment them with automation. But with all this data – and dramatic improvements in computing power – we have a phenomenal opportunity to do and achieve even more."

Robert Cardillo,
Director – NGA

2017

"The GEOINT discipline has **grown beyond the limits of human interpretation and explanation**. The explosion of available data diminishes the comparative advantage collection provides. Instead, automated processing, advancing tradecraft, human-machine collaboration, and the ability to anticipate behaviors will provide us a new advantage."

Robert Cardillo,
Director of NGA

A call to action: the urgency behind the DoD's adoption of AI

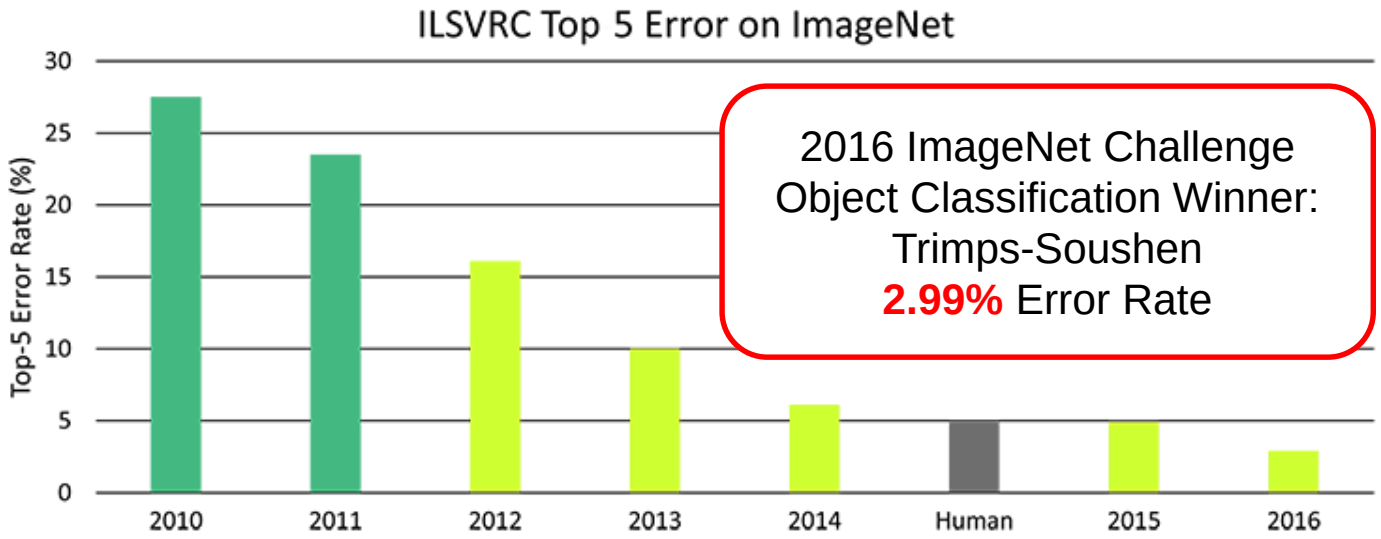

2016 ImageNet Challenge
Object Classification Winner:
Trimps-Soushen
2.99% Error Rate

Graph from https://www.dsiac.org

14M training images
1,000 object categories

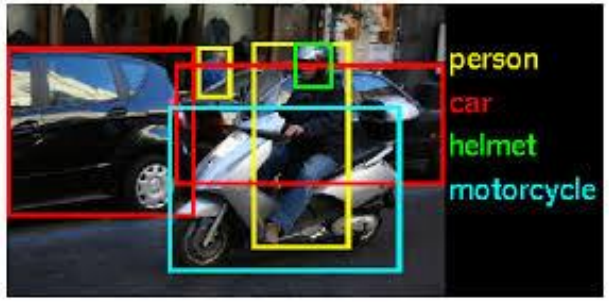

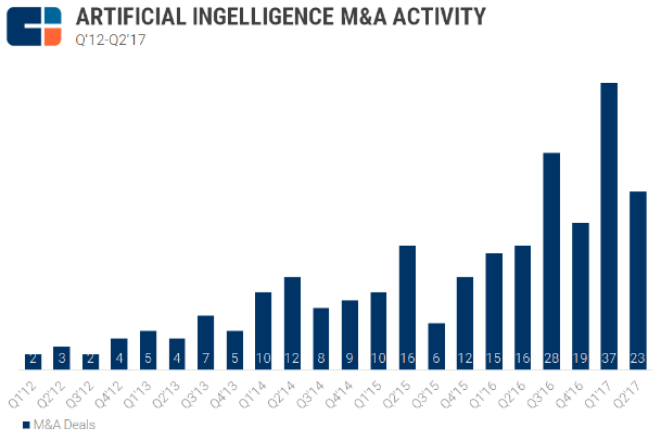

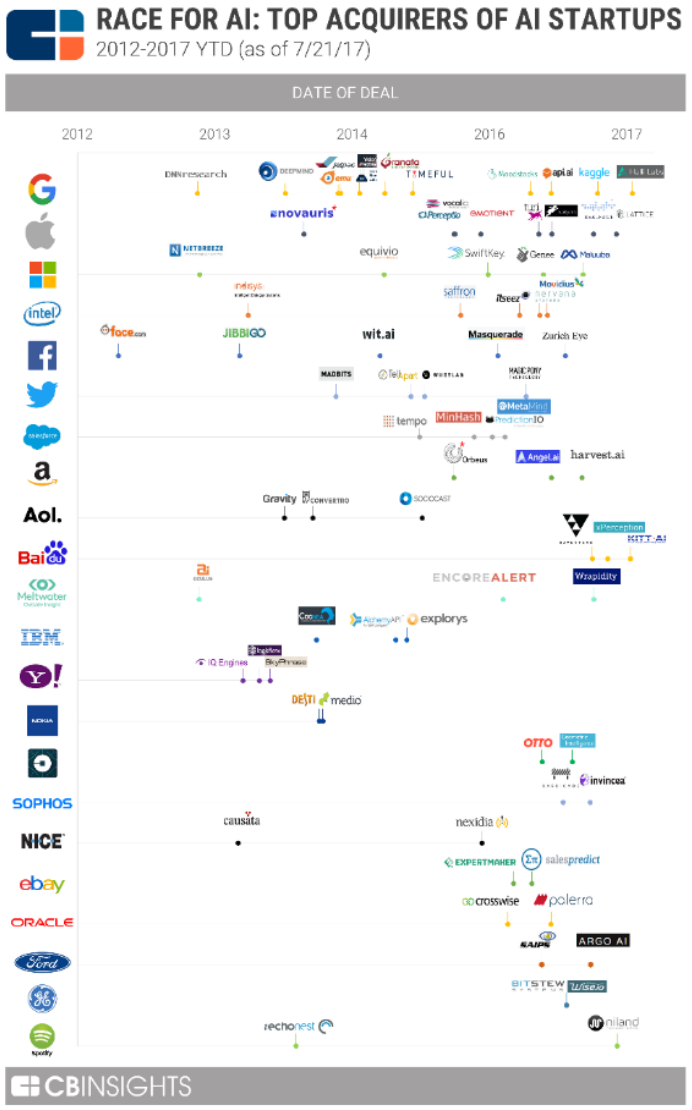

A call to action: the urgency behind the DoD's adoption of AI

"DoD has yet to embrace the transformational capabilities of artificial intelligence (AI) and machine learning (ML) across the Department. Together, they have the power to impact every corner of DoD, including force protection, training, logistics, recruiting, healthcare, C4ISR, cyber operations, and more. The significance of AI and ML is akin to the first and second offsets that took advantage of nuclear weapons and precision munitions and stealth technology, respectively. Underscoring the importance of applications of AI and ML across the Department is critical to create and sustain the asymmetric advantage required to outpace our adversaries."

- Defense Innovation Board, 2017

Agenda

Harris Corporation Introduction

A Call to Action: The Urgency Behind the DoD's Adoption of AI

Deep Learning Research at Harris

Harris's Work to Scale Deep Learning for Defense

Q&A

Harris deep learning research and development

Harris Corp. has been internally investing in taking state-of-the-art deep learning technologies and applying them to remote sensing and geospatial intelligence customer problems

How long?

- Harris has been working on deep learning for over five years

How much?

- Multimillion dollar internal research and development investment in the last three years
- Additional commercialization investment

Investment Approach

- Research
- Pilot Projects
- Software tool development
- Commercialization

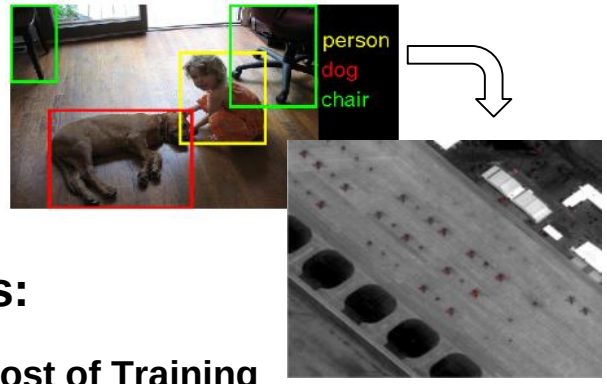

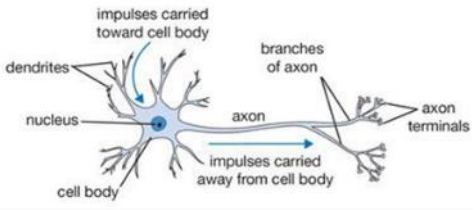

Focus Areas:

Reducing the Cost of Training

- Reduce dollar cost, human cost, and computer cost of building new models

Extensibility

- Ability to quickly redeploy and repackage tools to support new problem sets

Support multiple types of data

- New sensors and data fusion

Automation

- Ability for processes to interface to tools, removing human from the loop

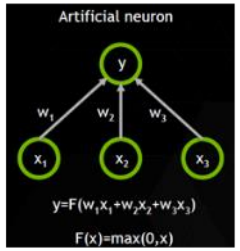

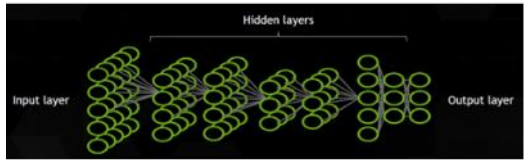

From NVidia

Harris deep learning R&D
Still imagery applications – automatic feature detection – 2D overhead

- Near ceiling performance in Pan, RGB, MSI
- Robust against occlusions, orientation, image quality

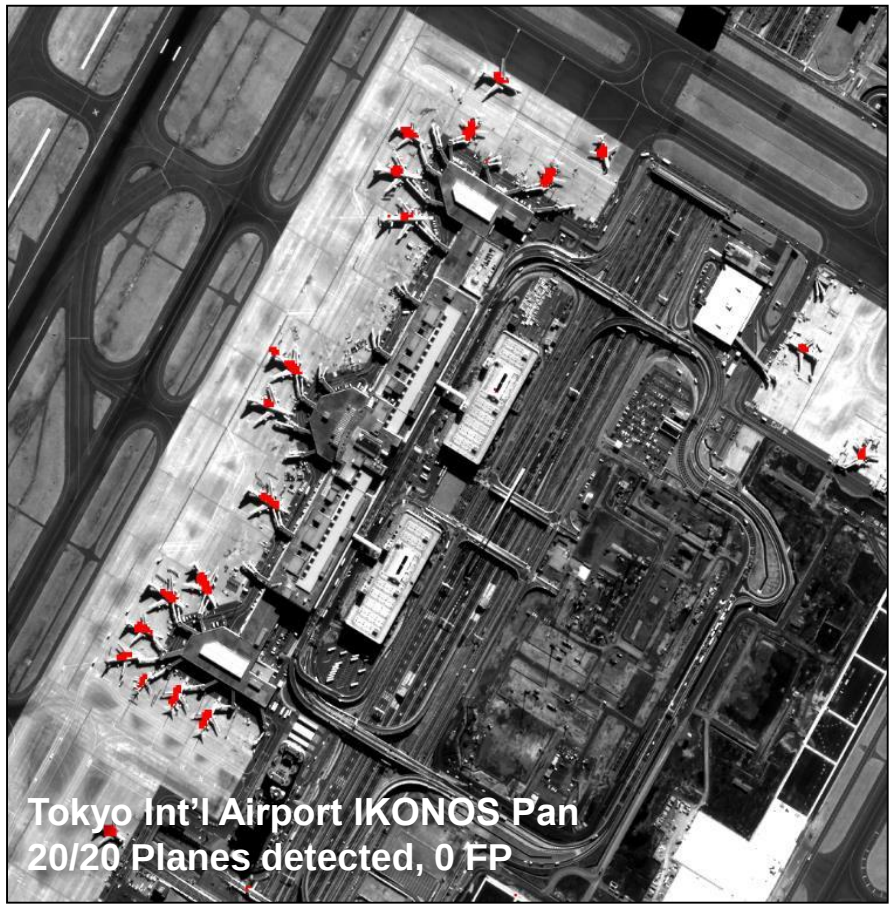

Sample Targets Tested:

- Airplanes
- Storage Tanks
- Sports Stadiums
- Athletic Fields
- Smokestacks
- Cooling Towers
- Clouds
- Crosswalks
- Swimming Pools
- Buildings
- Paved Roads
- Overpasses/Cloverleafs
- Tollbooths
- Beaches
- Cemeteries
- Wind Turbines
- Orchards & Row Crops

Harris deep learning R&D
still imagery applications – wind turbine blade inspection – 2D natural

Which wind turbines need repair in this wind farm?

BladeEdge℠
an EdgeData Company

http://edgedata.net/bladeedge/

- Drone-captured data
- Thousands of raw images organized
- Deep learning analyzed
- Damage identified and categorized
- Prioritize repair needs and budget
- Change detection through time
- Turbine life-cycle management

Harris deep learning R&D
still imagery applications – wind turbine blade inspection – 2D natural

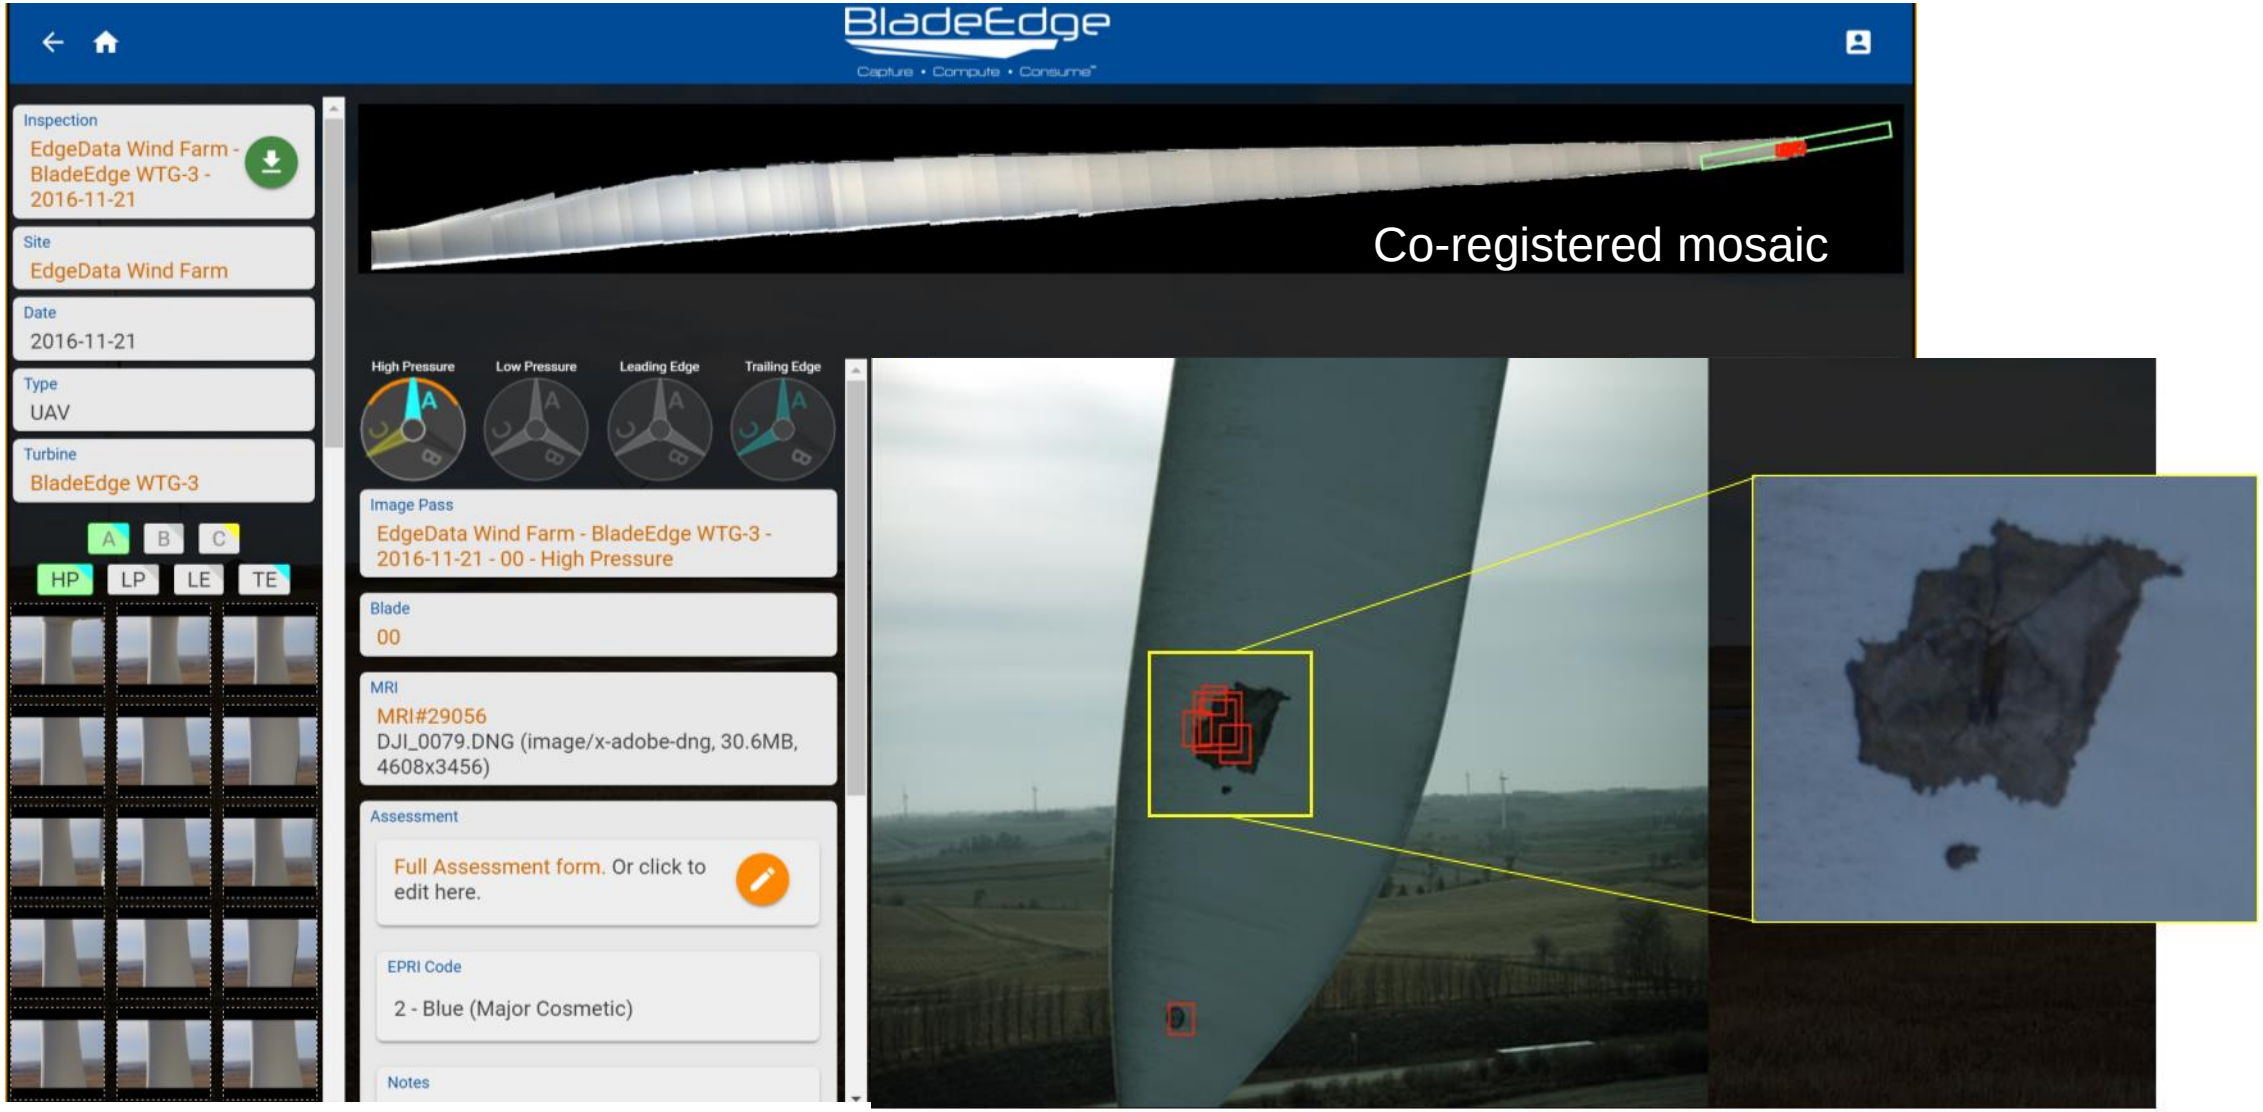

Harris deep learning R&D
motion imagery applications – traffic cams and road conditions

- Real-time ground weather intelligence system
- Data from traffic cam videos

- Scene detection instead of target detection
  Is it raining? Are the roads wet? Is snow present?

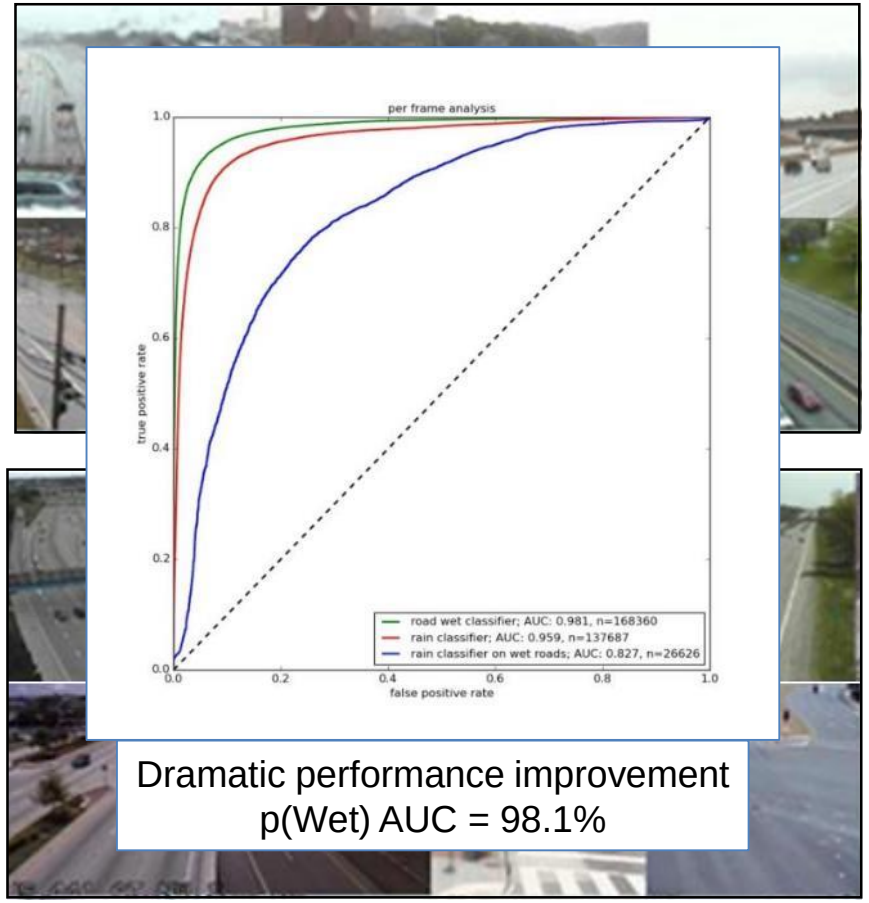

Dramatic performance improvement
p(Wet) AUC = 98.1%

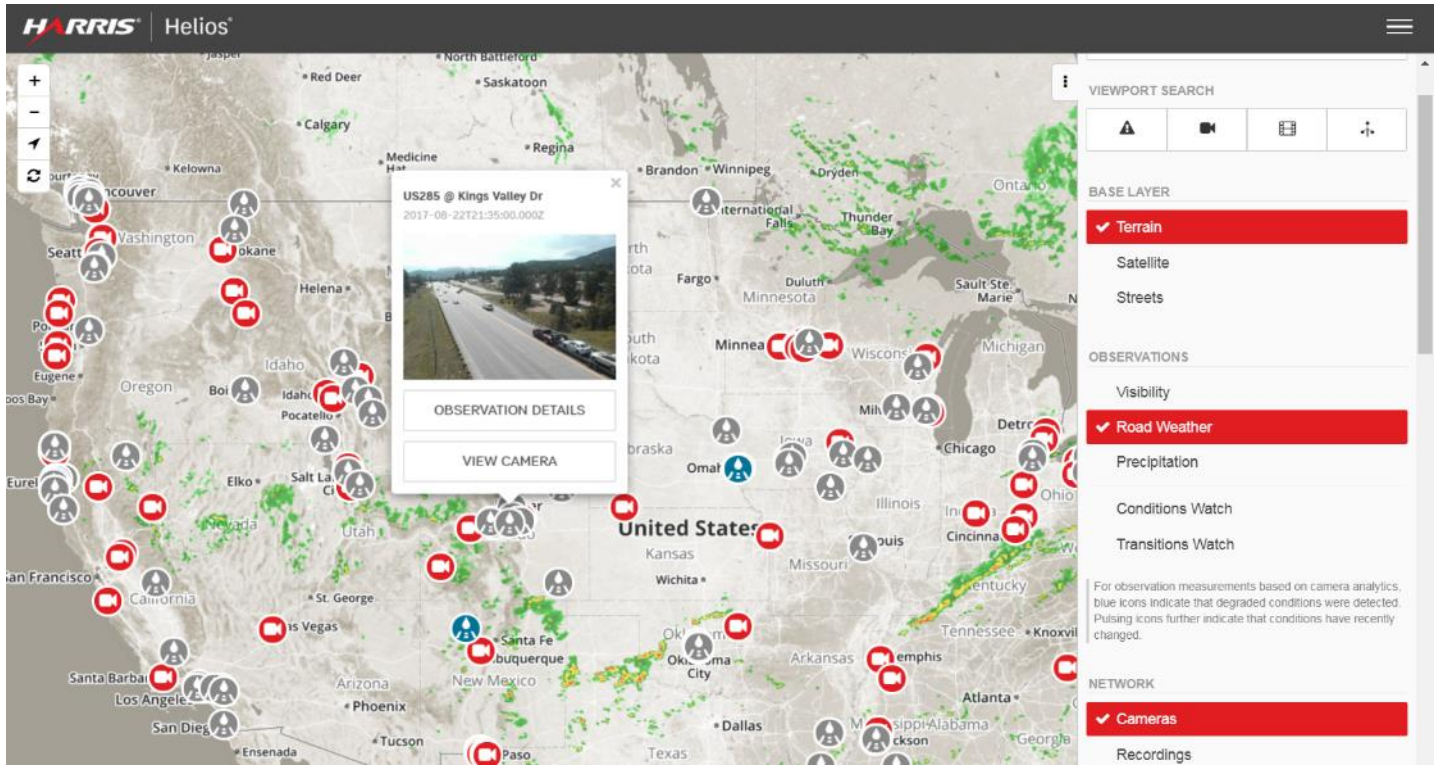

In use today in operational commercial product
https://www.helios.earth/

Harris deep learning R&D
motion imagery applications – Helios® for autonomous vehicles

Taking Harris Analytics to Connected Cars at the Consumer Electronics Show

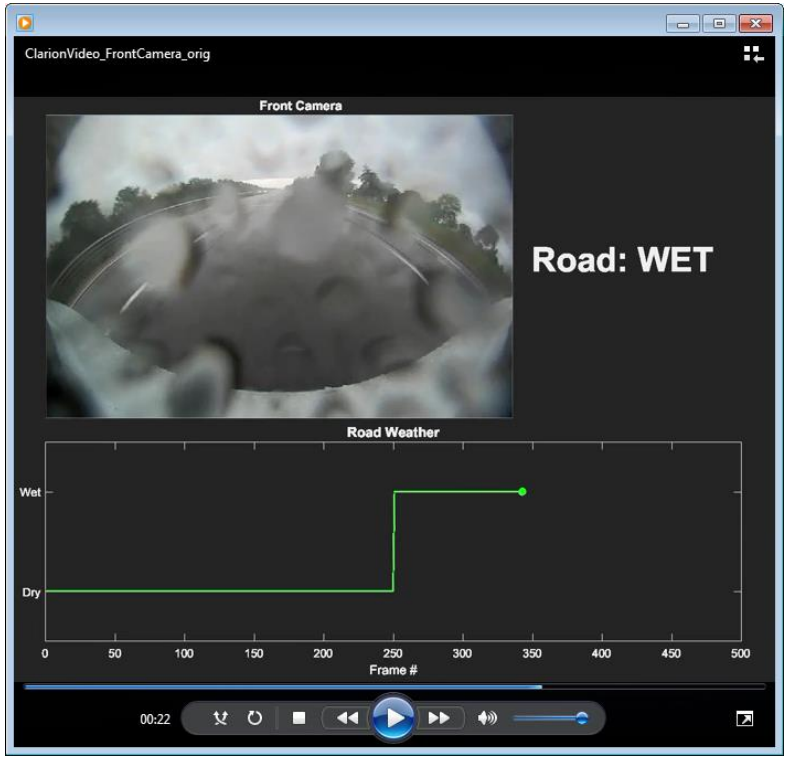

Harris deep learning R&D
motion imagery applications – high revisit rate still imagery

Problem: Clandestine airfields in South American countries used for illegal narcotic trafficking

Goal: Detect new airfields and determine activity levels with high temporal resolution data (Planet)

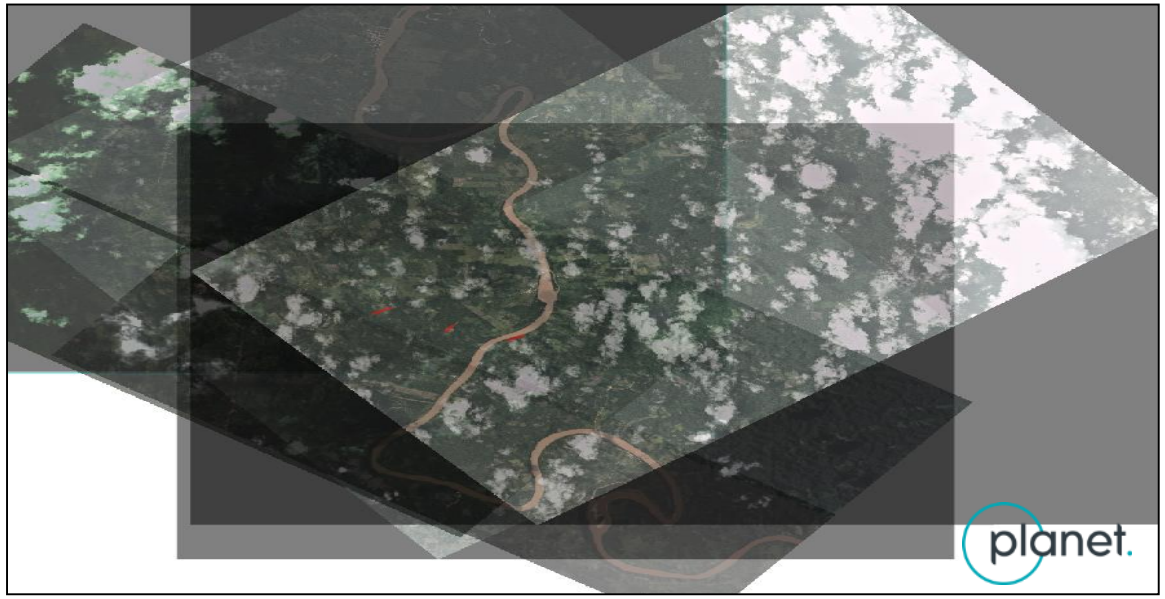

Harris deep learning R&D
motion imagery applications – high revisit rate still imagery

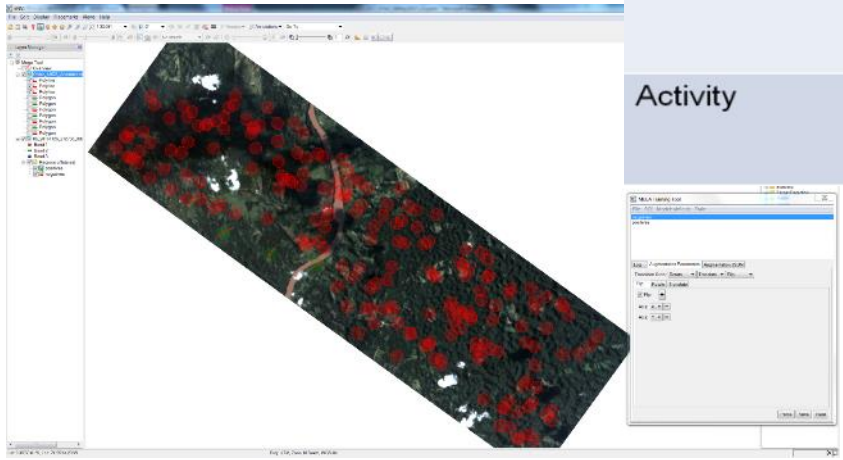

Date	Jan 3 '17
Scene	
Annotation	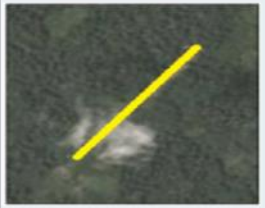
Activity	Airfield "Cratered"

AUC

	Single Frame	Temporal	LSTM
Fold 0	0.9555	0.9908	0.9887
Fold 1	0.9526	0.9927	0.9862
Fold 2	0.9666	0.9775	0.9838
Mean	**0.9582**	**0.9870**	**0.9862**

Accuracy

	Single Frame	Temporal	LSTM
Fold 0	91.84%	94.83%	94.61%
Fold 1	89.45%	95.95%	95.16%
Fold 2	90.53%	92.71%	94.35%
Mean	**90.61%**	**94.50%**	**94.71%**

Sources:

- Planet imagery
- DigitalGlobe EGD imagery
- Ecuadorean geoportal shapefile of known remote landing strips
- Google Earth imagery
- Wikimapia
- OpenStreetMap

Harris deep learning R&D
motion imagery applications – natural video action recognition

UCF-101 benchmark dataset

- 101 classes
- 13k video clips
- Approximately 30 hours of video

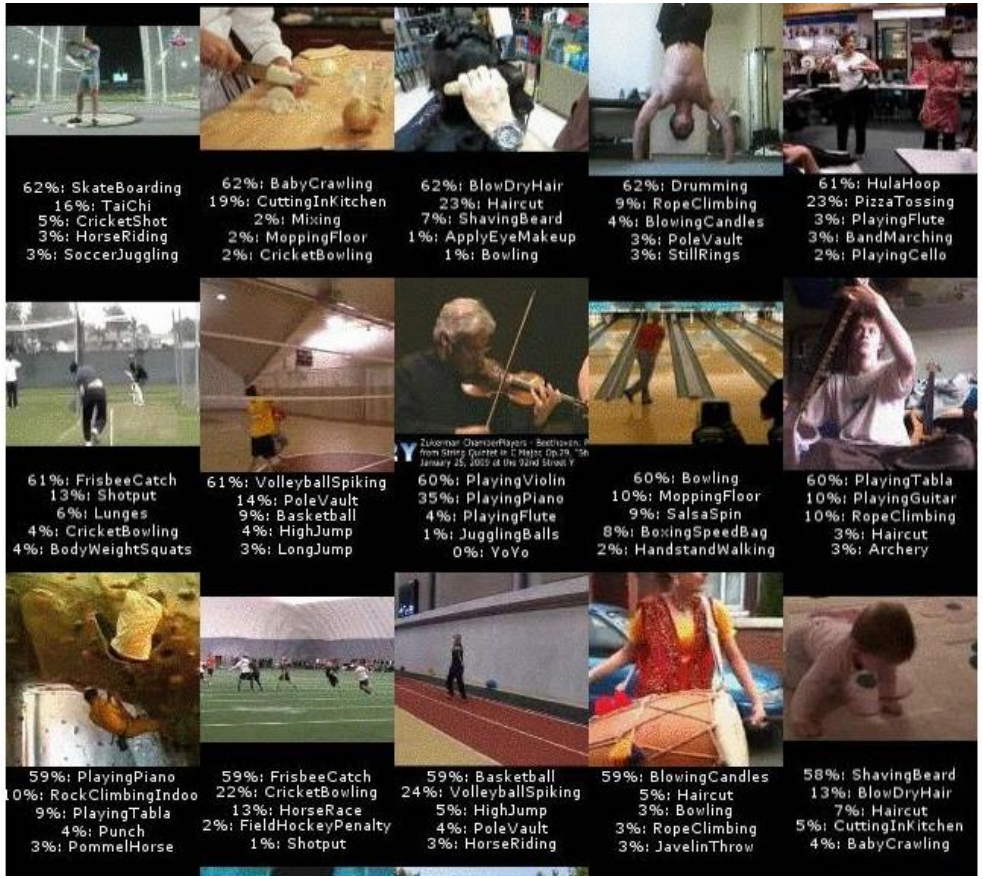

Among the most challenging open problems

- Demo of successful training of recurrent NN (LSTM)
- Model was learning despite severe reduction in data (length of clips and resolution of frames)
- Analogue for surveillance cameras or drone video
- Many potential applications for RNNs
  - Any problem that can be restated in terms of a sequence

Harris deep learning R&D
motion imagery applications – overhead video analytics

Transition from forward-facing security camera type video to offset overhead video

- Fewer pixels on target
- Less stable
- Changing orientations, modalities, zoom levels

Source: http://www.newsweek.com/2014/04/25/biometric-surveillance-means-someone-always-watching-248161.html
Credit: JOHN MOORE/GETTY

Source: http://www.viratdata.org

Harris deep learning R&D
LiDAR applications – introduction

LiDAR sensors have very powerful discriminatory capabilities. Can DL models learn to recognize objects and terrain types?

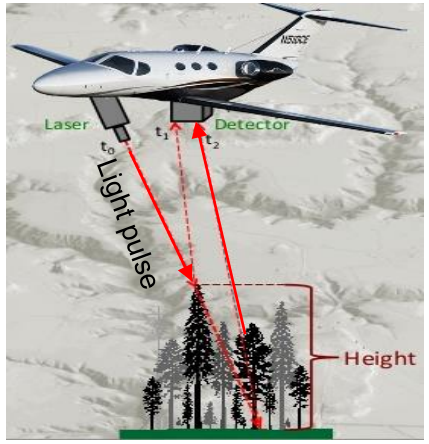

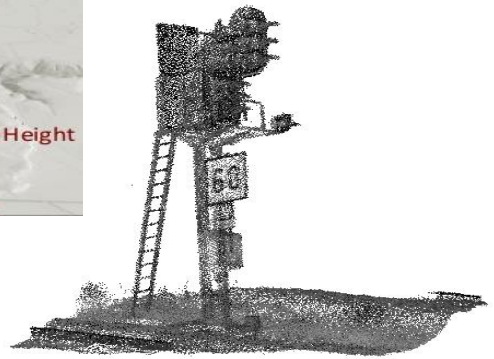

Many Challenges

- Raw 3D point clouds are difficult to use compared to images
- Typically very large and spatially dense, yet sparse in Z-dim
- Complementary to imagery, could be used together if available
- Harder to label
  - By individual point or 3D bounding box
  - Software tools less developed (or don't exist)
- Not much deep learning research on airborne LiDAR classification

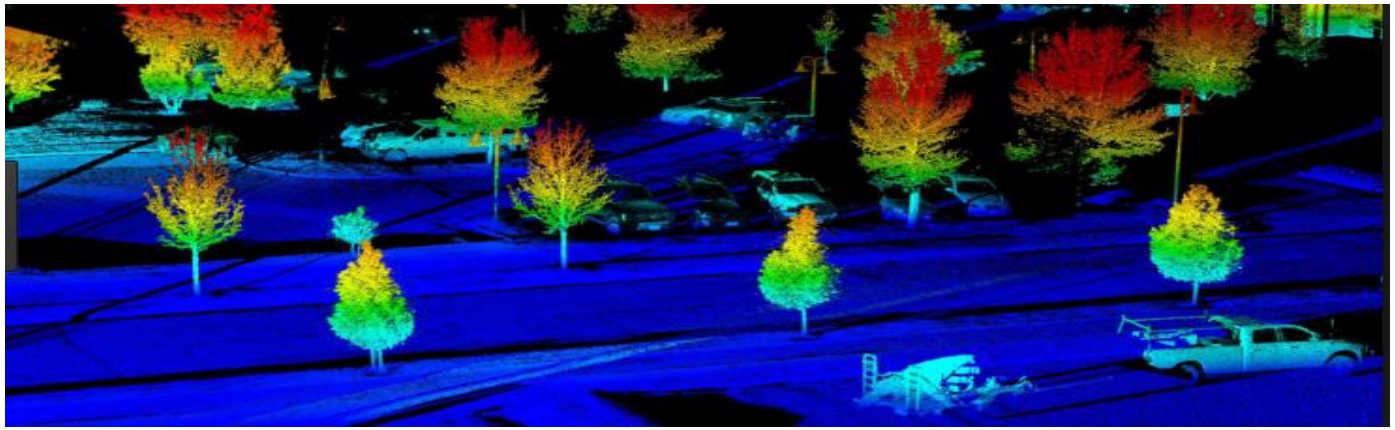

3D point cloud color-coded by height

End-to-End Approach for Success

- Understanding phenomenology from sensor to analyst allows for targeted application of deep learning
  - Deep learning has been adopted at multiple points in production workflow supporting automation and product generation

Harris deep learning R&D
LiDAR applications – summary

Object Detection in LiDAR Point Clouds

- French National Railroad (SNCF) asset inventory project
- Extension of 2D ConvNet to 3D data source
- Preliminary results are very encouraging
- Finds variety of 3D objects (signals, crossings, boxes, poles)

Example Target: Railroad Signals

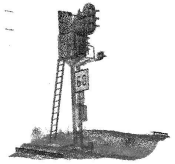

Heat Map (all detections shown)

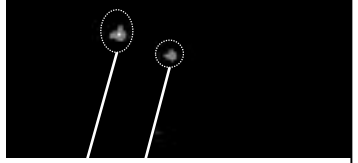

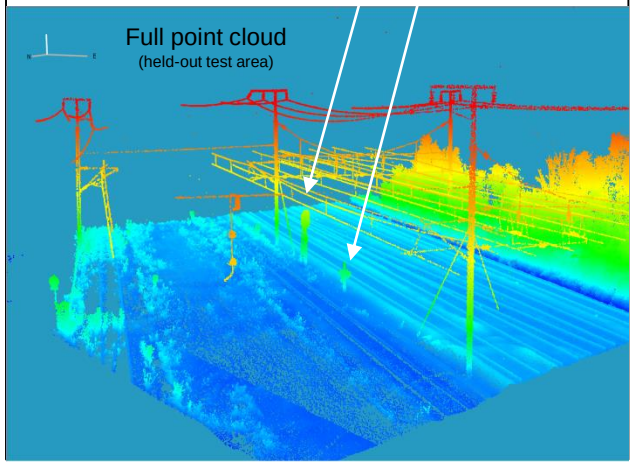

Geiger-Mode LiDAR Bare Earth Point Labeling

- Proof-of-concept for point labeling in Geiger-Mode data
- Limited training data w/significant errors
- Discovered effective signals in data by constructing voxel grids with
  - mean/min/max/stdev of elevation and intensity
  - Multiple neighborhoods of gridded data used as inputs for CNN
  - Post-processing with CRF to enforce nearest neighbor dependence

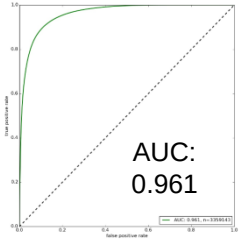

Bare Earth Probability Map

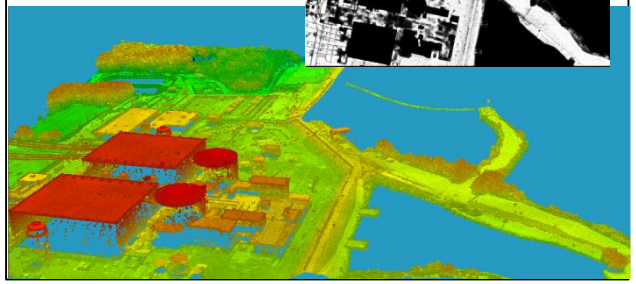

Point Clouds from Synthetic Image Pairs

- Point cloud data derived from stereo pairs of imagery
- MEGA analytics agnostic to multiple data generation approaches
- Contains information about multiple collects and sensor model

Generated point cloud from MSI/HSI

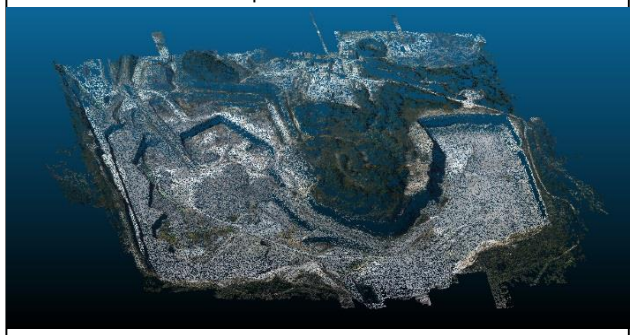

Model must understand context of “missing” data

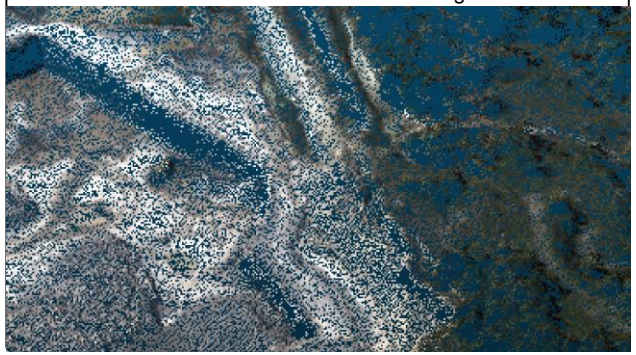

Harris deep learning R&D
Applications lessons learned and takeaways

Pros:

Now possible to build detection algorithms with minimal prior knowledge of the target and minimal image science background

High 90% accuracies achievable by amateurs with sufficient label data and using appropriate neural network architecture

GPUs allow neural net training and inference times to be timely (minutes instead of days)

Agenda

Harris Corporation Introduction

A Call to Action: The Urgency Behind the DoD's Adoption of AI

Deep Learning Research at Harris

Harris's Work to Scale Deep Learning for Defense

Q&A

The art of scaling ML

Problem → Machine Learning → Answers!

Mobile Access

Understand problem

PhD

Training Data

Hardware

Data ingest

Present solution

Classification Export Controls

Labels

Curation

Lots of hardware

Sensor knowledge

Retrain model

Performance monitoring

Scheduling

Geographic Metadata

Orthorectification

Refinement

Notifications

Prioritization

Breaking it down

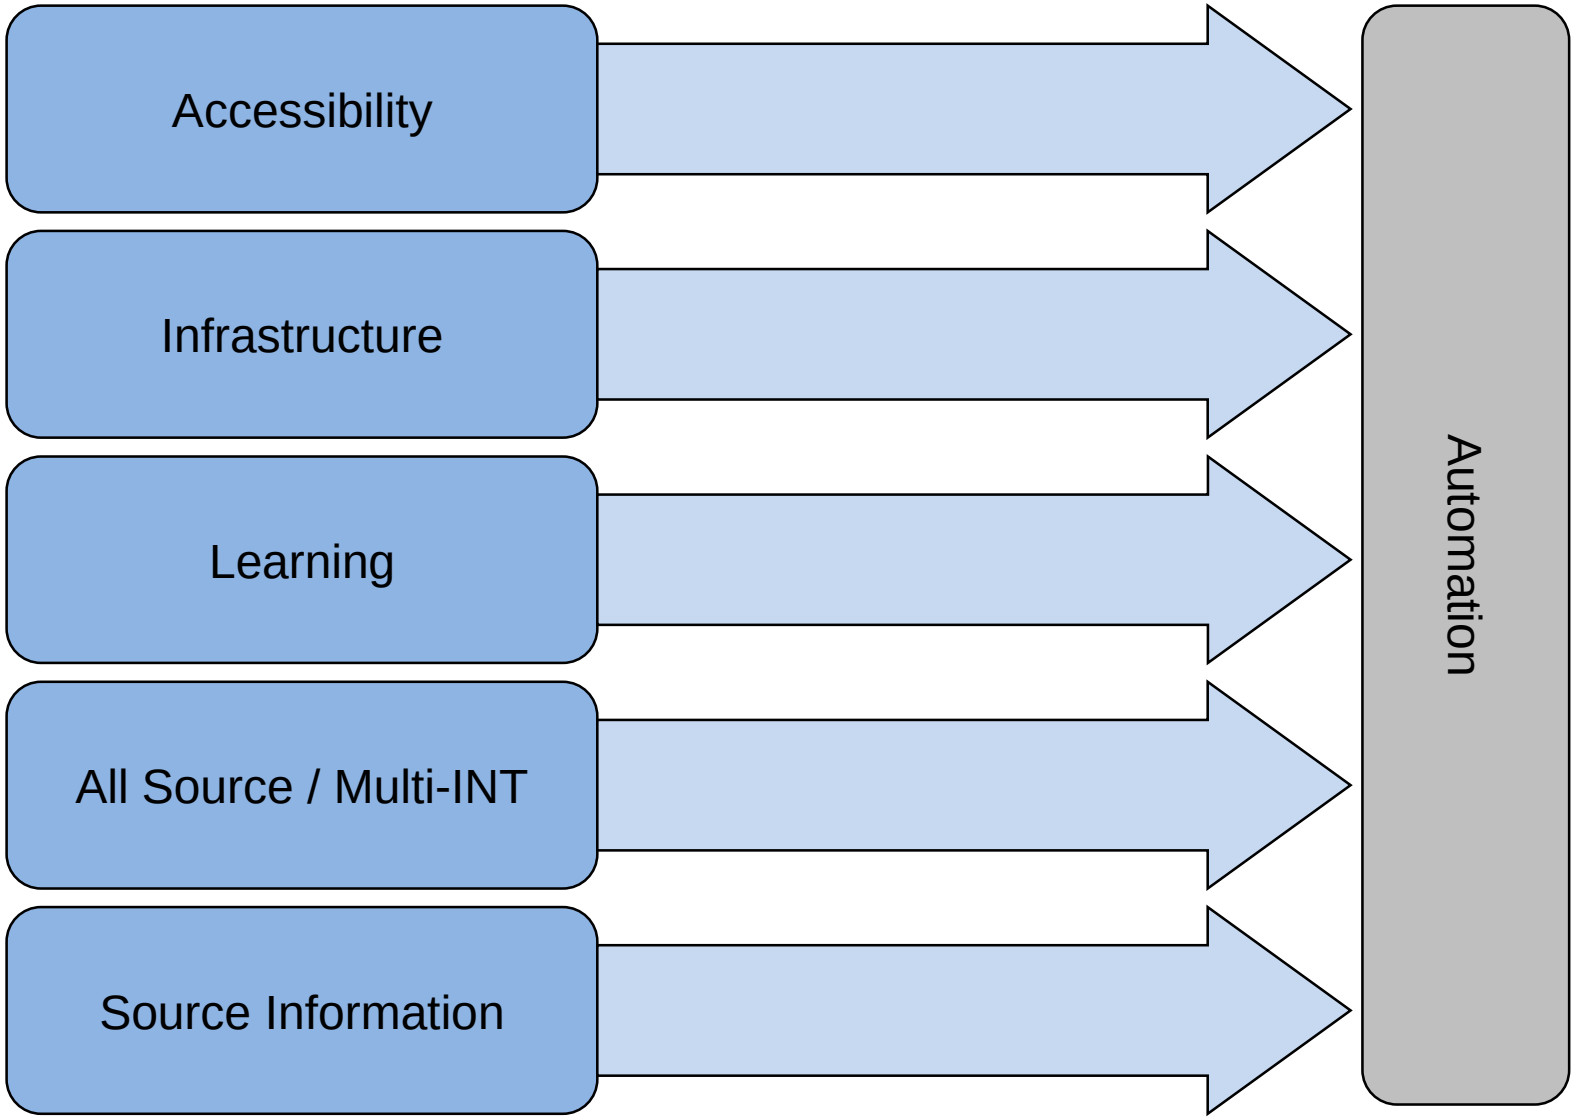

Rinse and repeat

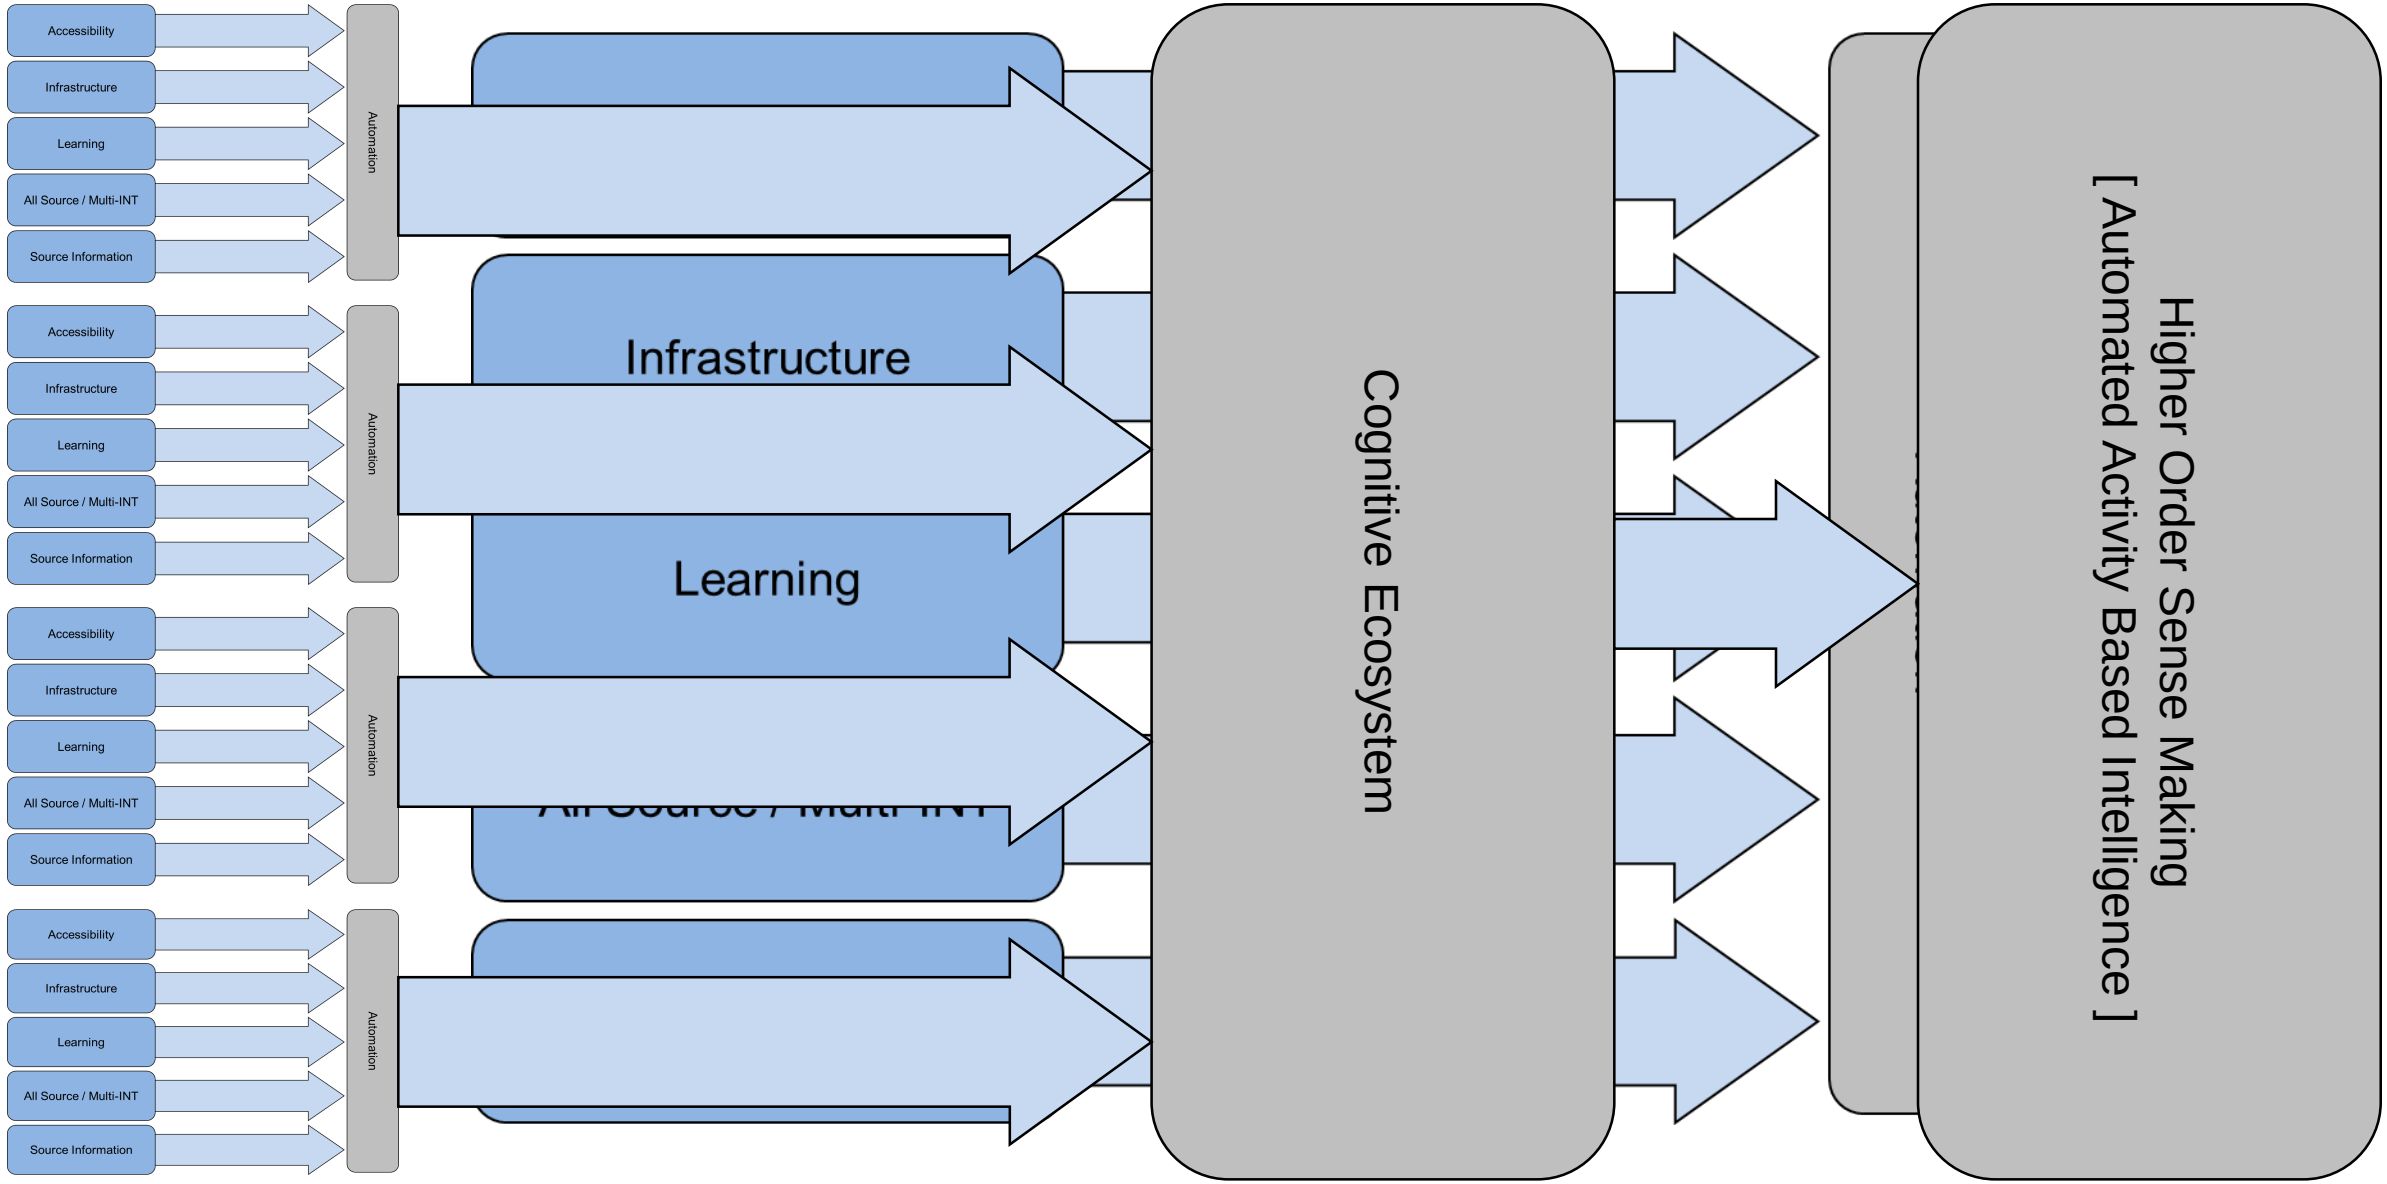

Harris' work to scale deep learning for defense source information – labeled data – the currency of ML

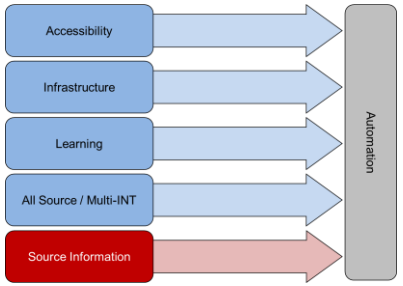

Real-world problems don't have an ImageNet for training.
How do we overcome the burden of labeled data?

Just Go Get It

- Manual harvesting
- Crowdsourcing
- Use of (limited) published datasets

Do More With Less

Supervised

- Pre-train the network and transfer learn where possible
- Enforce domain adaptation if necessary
- Use synthetic data from models or simulators
- Use curriculum learning and dynamically balanced datasets

Unsupervised and Active Learning

- Use a deep embedding to optimize clustering of unlabeled data
- Iterate training of weakly supervised models to progressively improve training set
- Use Convolutional Autoencoders or GANs to generate embedding from unlabeled data
- Use Siamese NN to find similar labels or extract an embedding

Harris' work to scale deep learning for defense source information – manual harvesting – individual tools & techniques

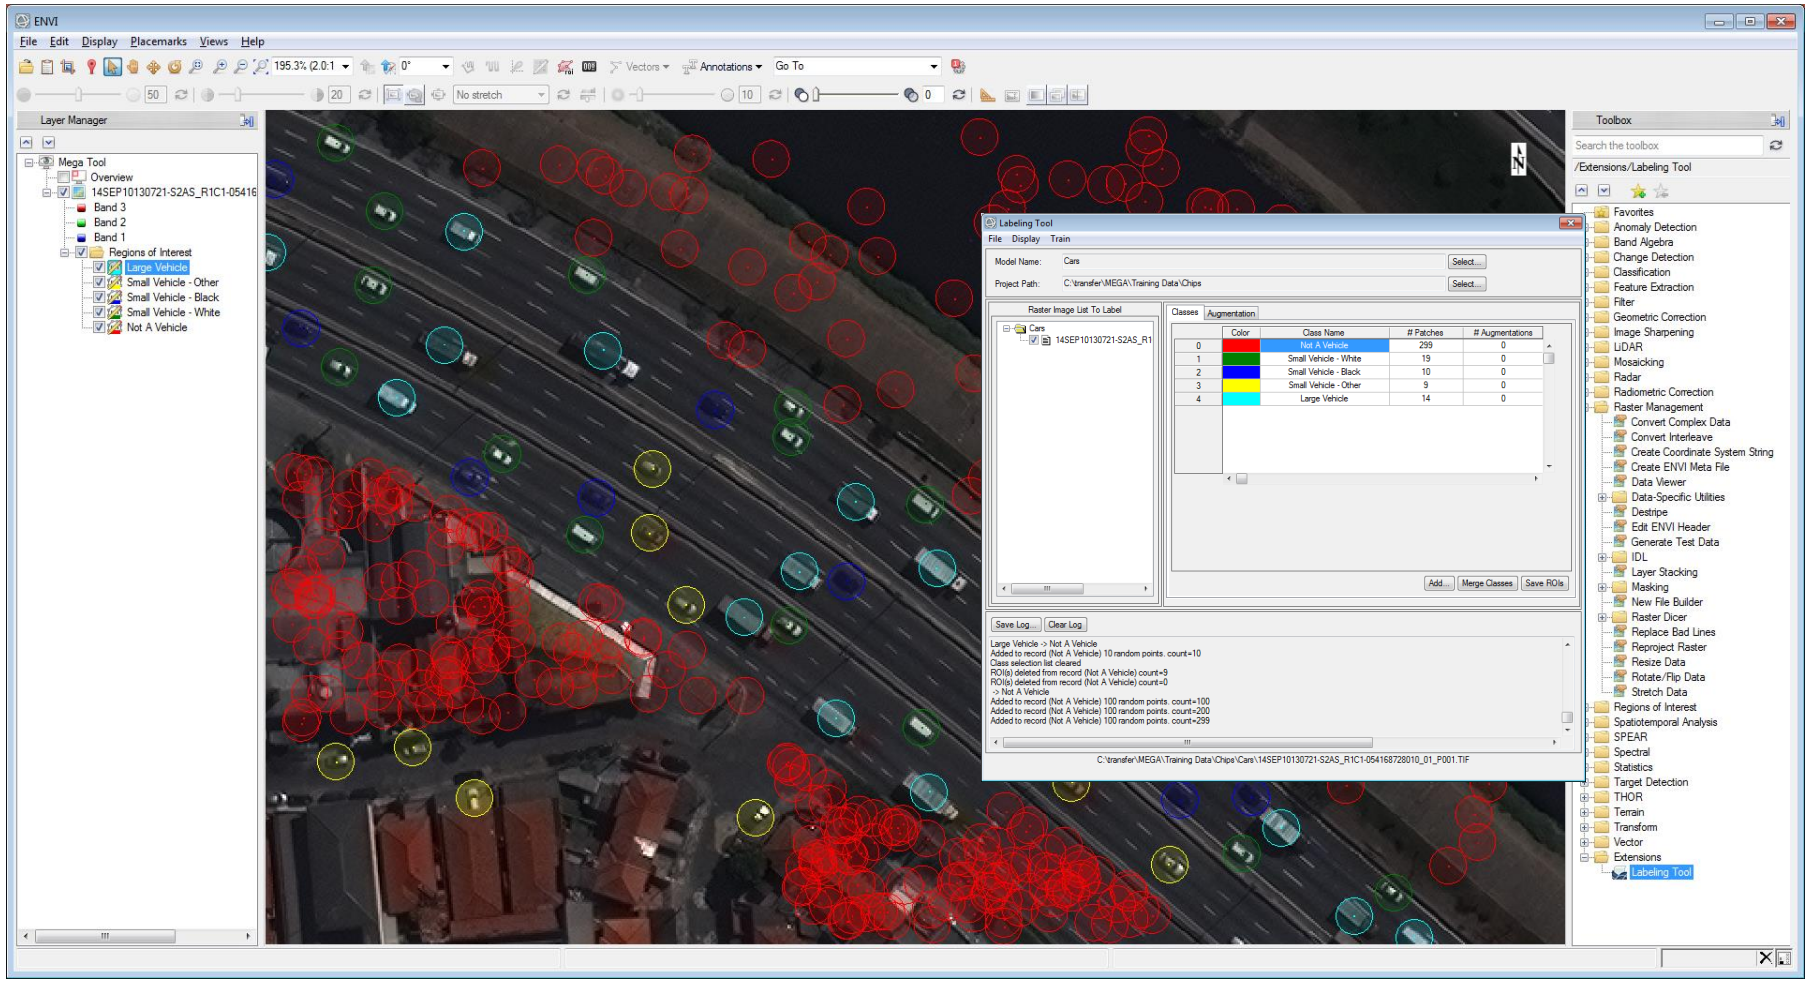

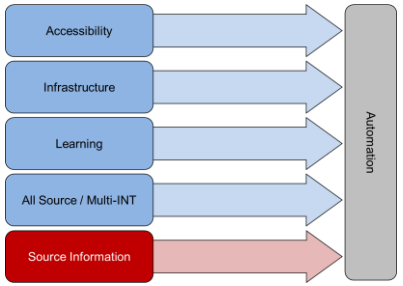

Harris’ work to scale deep learning for defense source information – manual harvesting – crowdsourcing

amazonmechanical turk
Artificial Artificial Intelligence

OpenStreetMap

taginfo

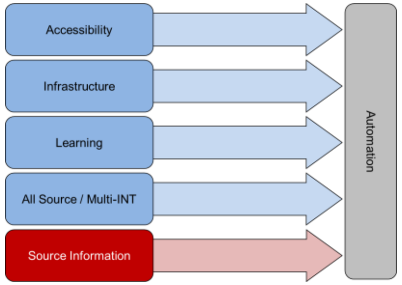
- Accessibility
- Infrastructure
- Learning
- All Source / Multi-INT
- Source Information

Automation

Exploit commercial tools and public-domain datasets

- Feature and point-of-interest information

Inexpensive and often fast

Need to mitigate Q/A requirements

- Some labeling error can actually be good during model training

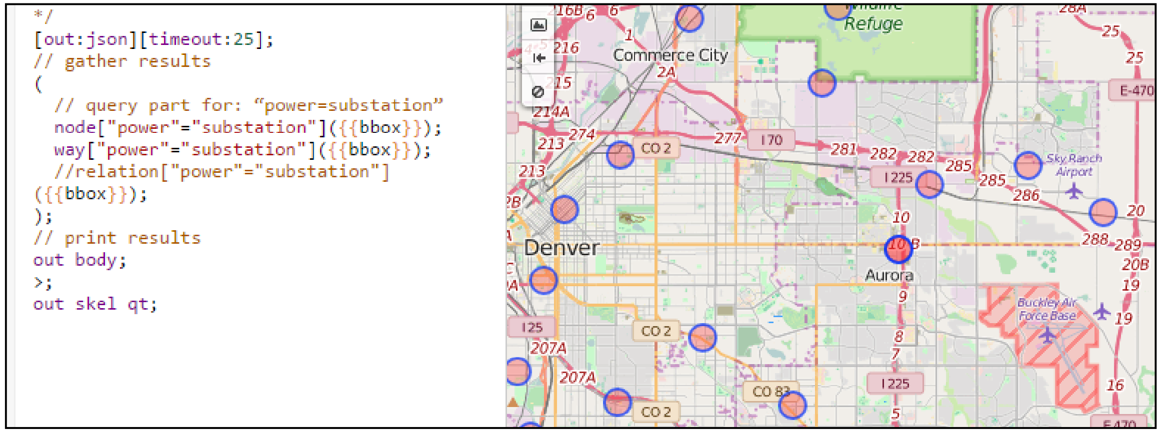

```
*/
[out:json][timeout:25];
// gather results
(
  // query part for: “power=substation”
  node["power"="substation"]({{bbox}});
  way["power"="substation"]({{bbox}});
  //relation["power"="substation"]
({{bbox}});
);
// print results
out body;
>;
out skel qt;
```

Harris' work to scale deep learning for defense source information – auto-clustering via deep embedding space

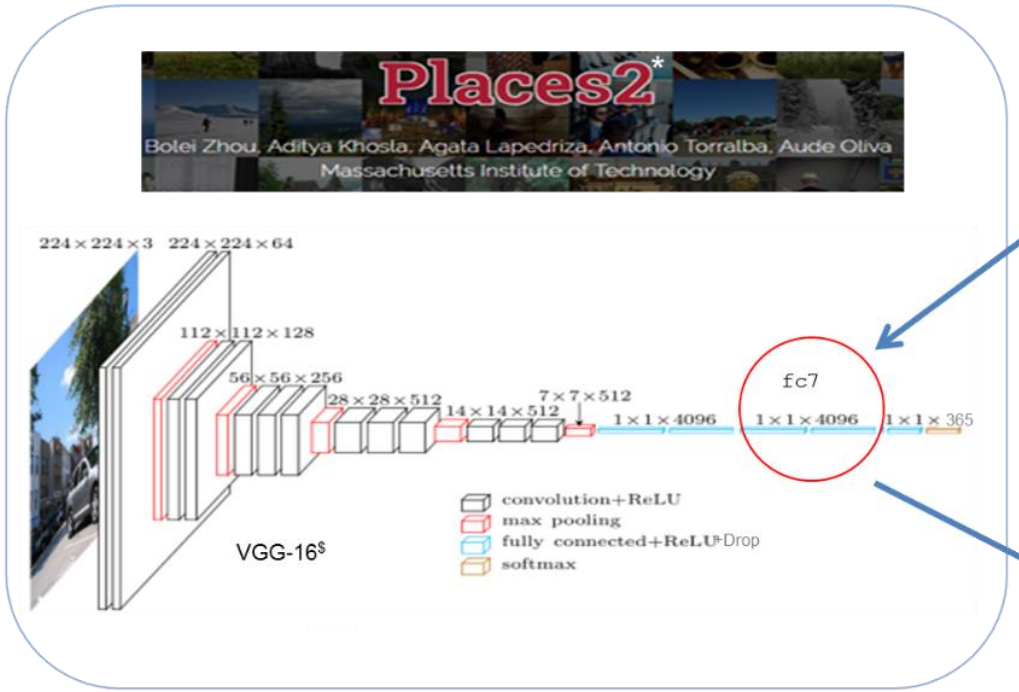

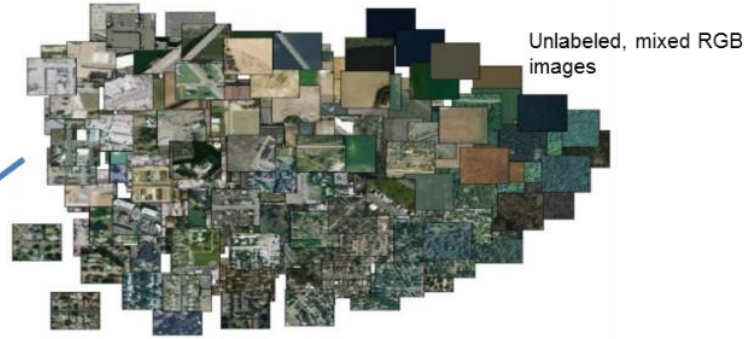

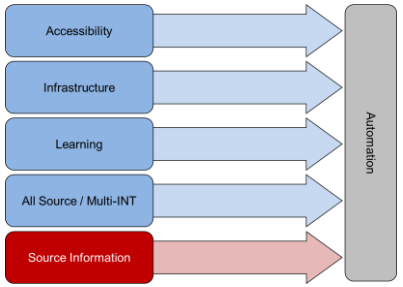

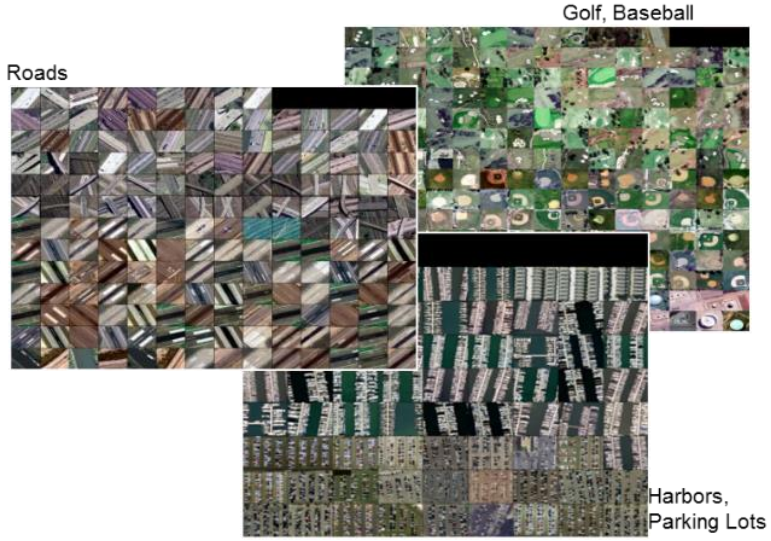

Deep Embedding from VGG-16 trained on Places365 scenes used to cluster mixed RGB imagery into like labels

* *Places: An Image Database for Deep Scene Understanding*, B. Zhou, A. Khosla, A. Lapedriza, A. Torralba and A. Oliva *arXiv:1610.02055*
§ Very Deep Convolutional Networks for Large-Scale Image Recognition K. Simonyan, A. Zisserman arXiv:1409.1556

Harris' work to scale deep learning for defense source information – synthetic label data generation

- 100% of training data synthesized using CAD models and Scene Simulator
- The trained model is applied to real imagery
- Successful detector produced for fighter jets in WV-2 Pan imagery
- Limiting factors: (1) content of scene generator and (2) quality of simulation

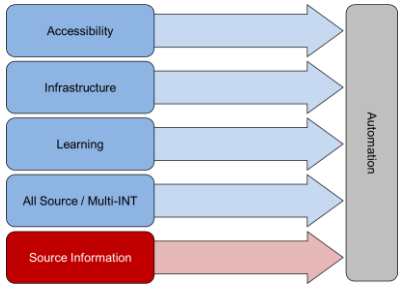

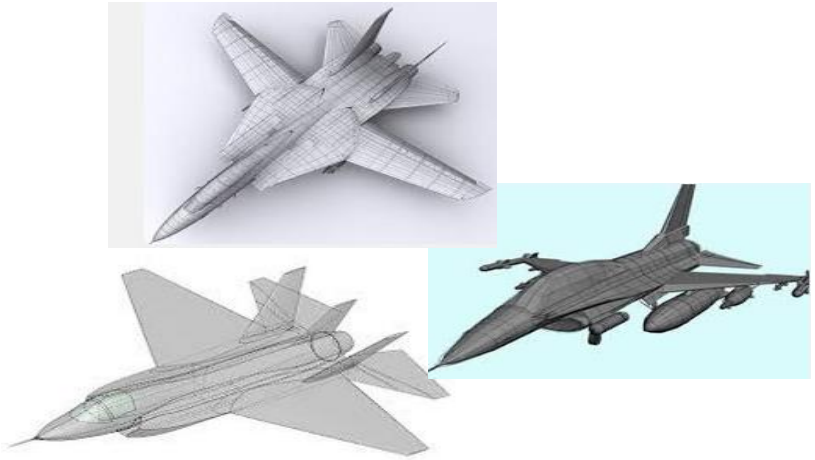

6 CAD models used

DIRSIG

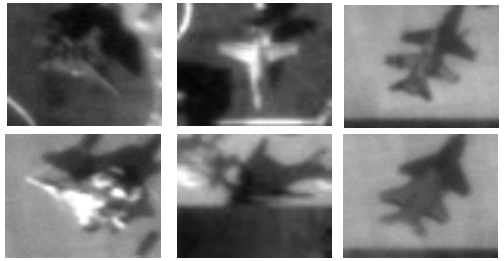

10 am, 7 degree look angle, Jan 1, Scene Azimuth 0

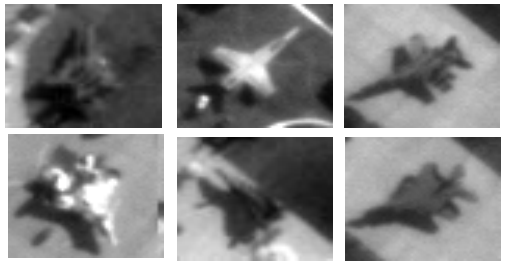

2 pm, 7 degree look angle, Jan 1, Scene Azimuth 225

Objects placed in scene at various geometries

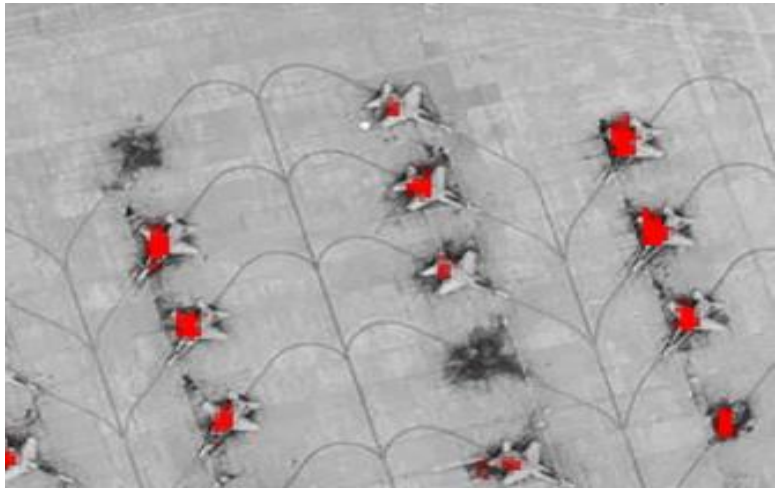

Heat Map for fighter jets in IKONOS Pan Imagery

Harris' work to scale deep learning for defense all source/multi-INT – holistic approach

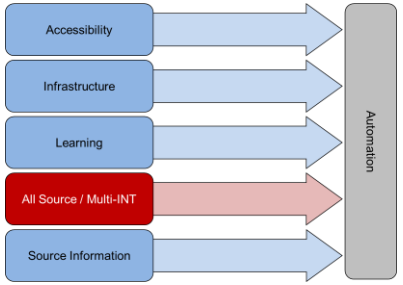

Context is critical in understand and deriving intelligence

Human analysts add context through traditional GIS approaches

- Shapes, points, annotations
- Machines can be taught to interpret data as stream to add human context or can be allowed to create their own

Goal is holistic view of world

Use what is appropriate to solve problem

- Keep all information available
- More good data and more good context allows enhanced intelligence and decision making

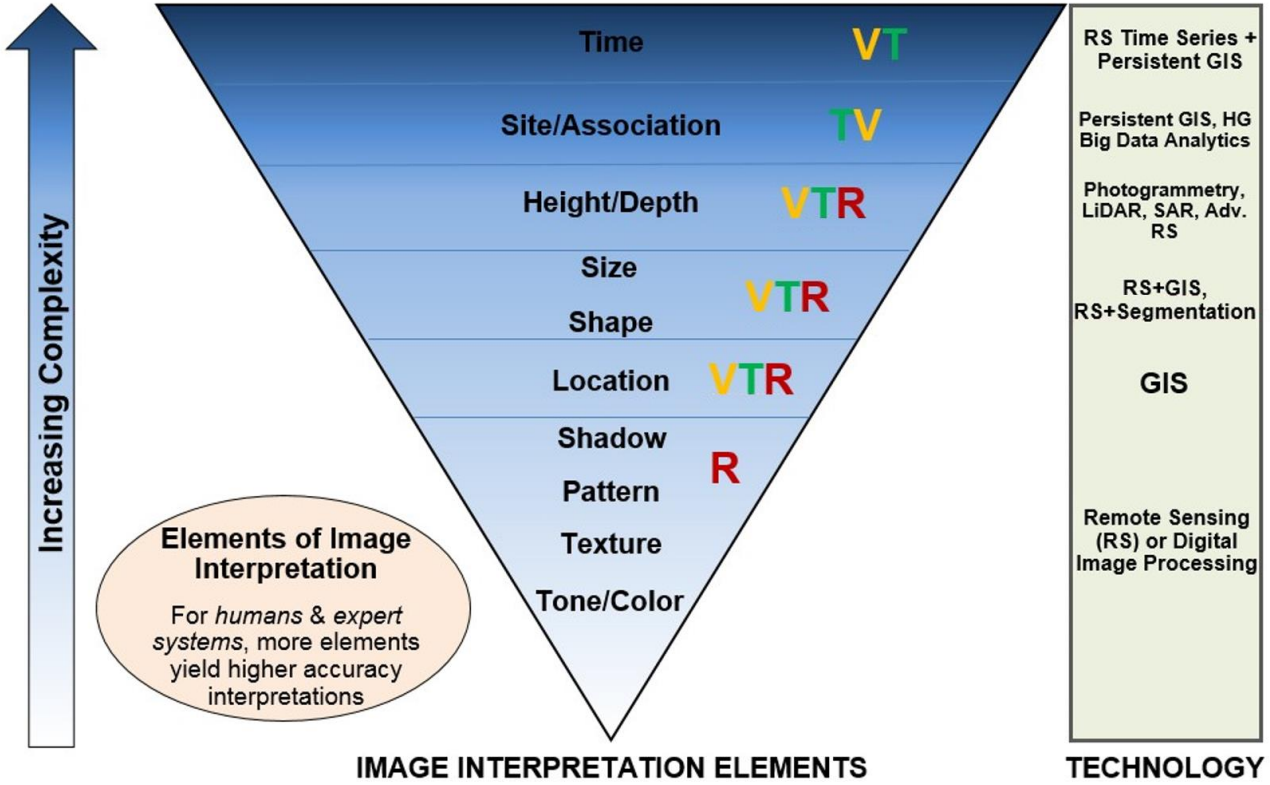

Harris' work to scale deep learning for defense
all source/multi-INT – multiple data inputs & conv features

What's the best way to handle multiple data inputs?

1. Different data dimensions
   input 1 = [200, 200, 3] Panchromatic image
   input 2 = [50, 50] Multispectral image

2. Different sensor modalities
   input 1 = Multispectral image
   input 2 = LiDAR points (or derived image product)
   input 3 = SAR image

3. Disparate data sources
   input 1 = natural image
   input 2 = audiogram
   input 3 = RADON transform

4. New data?

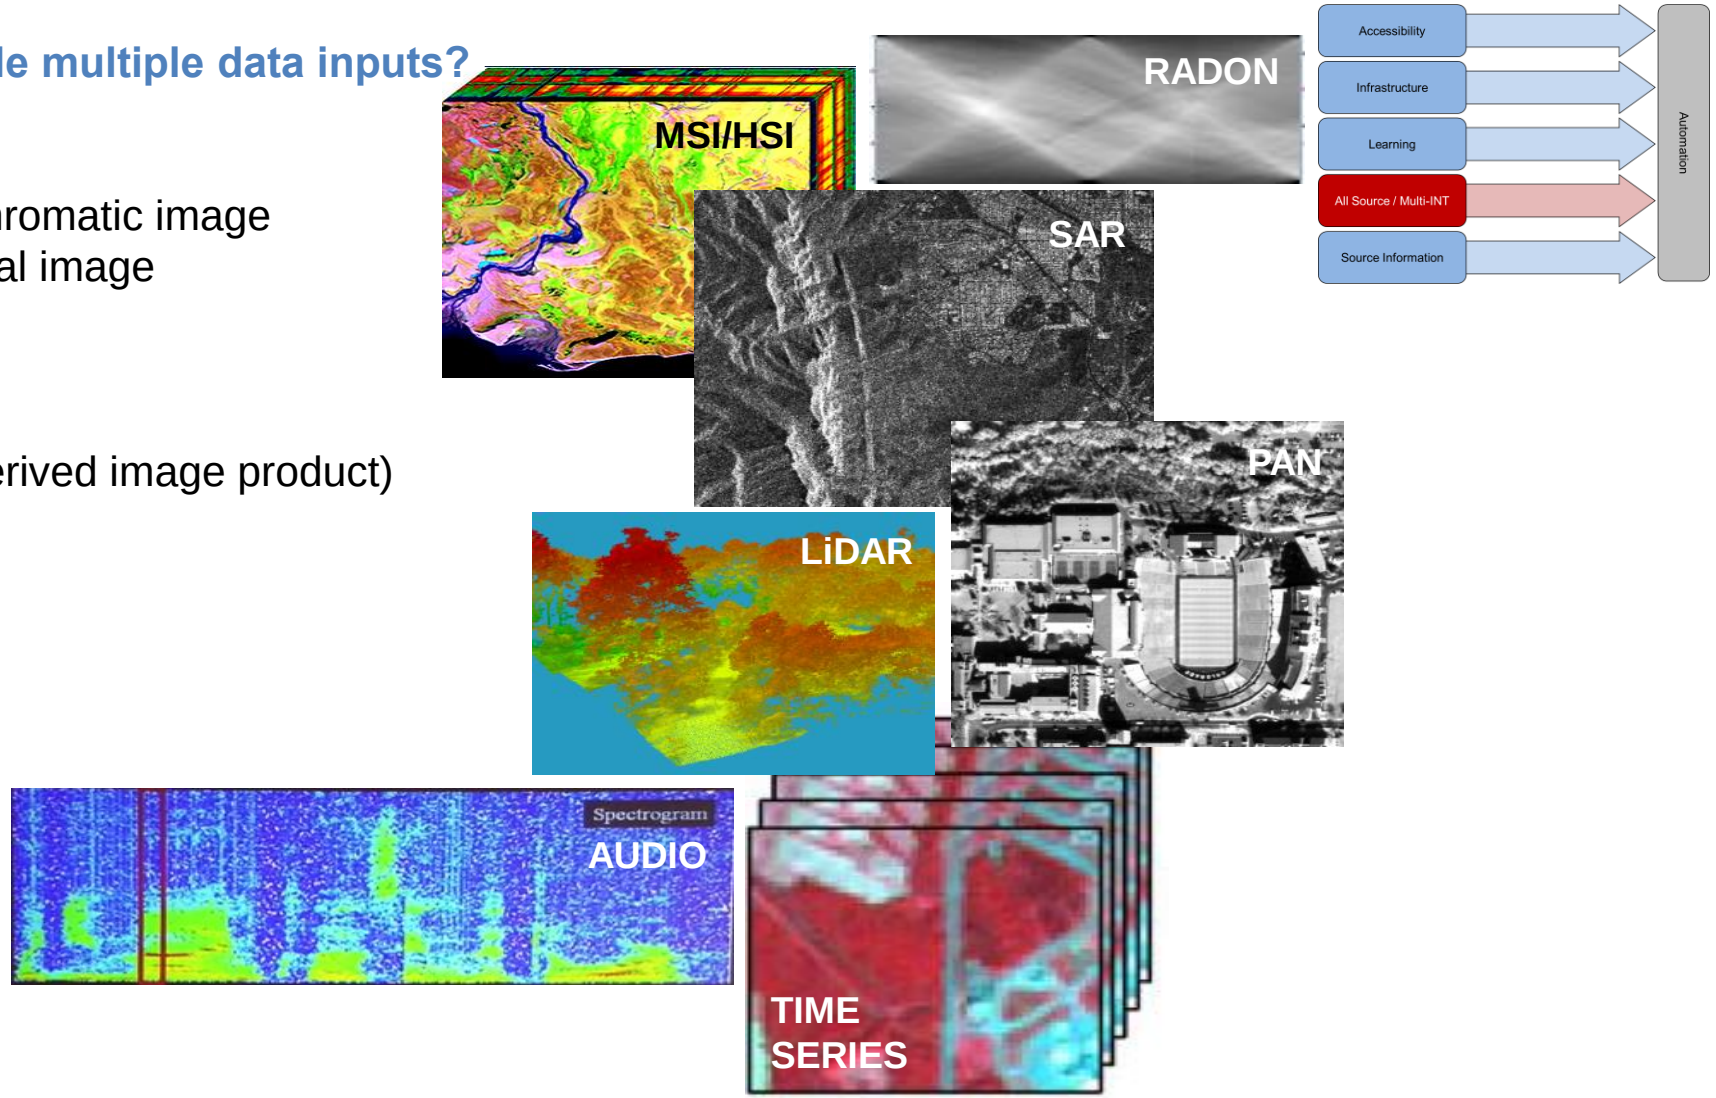

Harris' work to scale deep learning for defense
all source/multi-INT – multiple data inputs & conv features

Supporting many data sources enables higher level decision making

- Structured / unstructured correlated data
- Multi-modalities / multi-format

DL support for multi-stream processing enables cross-functional collaboration

- Bridges gap between agency, function, mission

Make the most of what you have – use all available context to derive answers

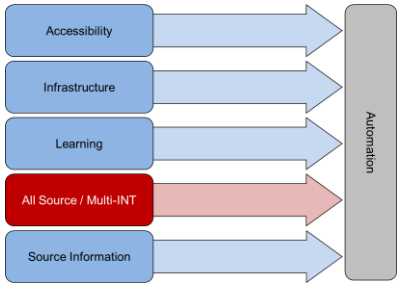

A very flexible model architecture

- Enables natural data fusion – all data input in original form
- Maximizes feature extraction from complimentary data independently
- Effective way of providing critical context for targeted task

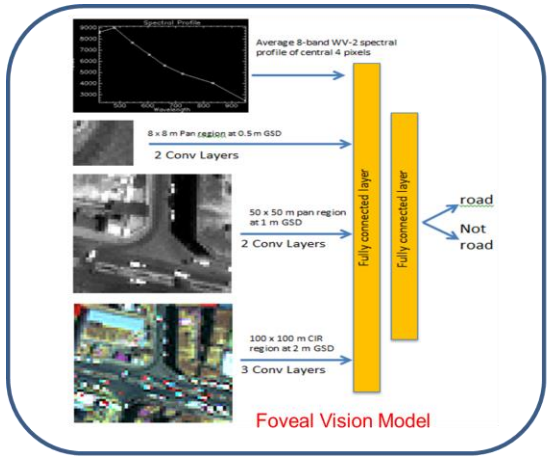

4-Input Data Streams

- Progressive scales of observation
- Precise co-registration not required
- Spectrum provided as separate input

Road networks in Mombassa

- R0: Detected roads overlaid on TrueColor RGB

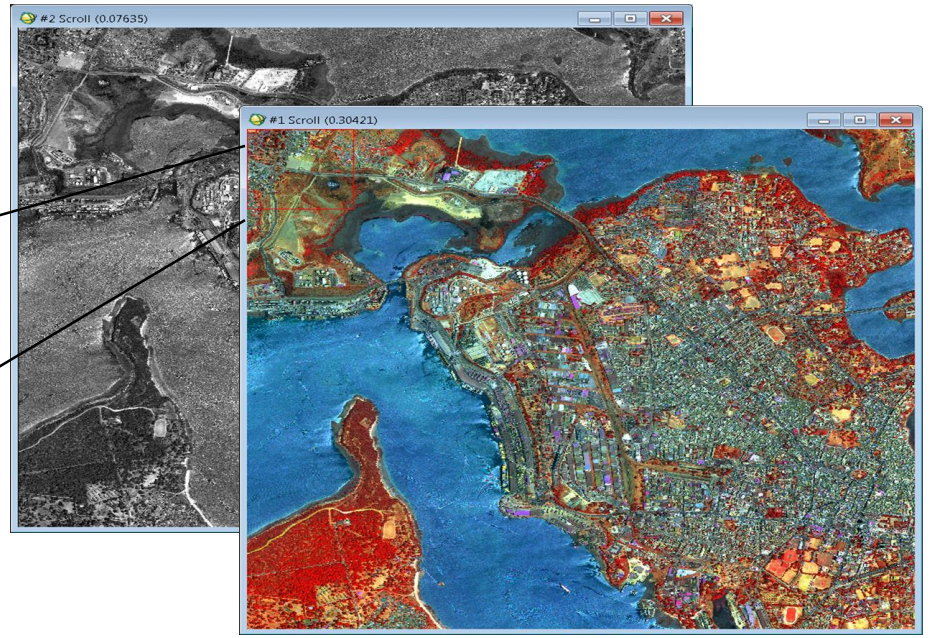

Harris' work to scale deep learning for defense
all source / multi-INT – heat map improvements with multi-stream

Multi-Stream Inputs

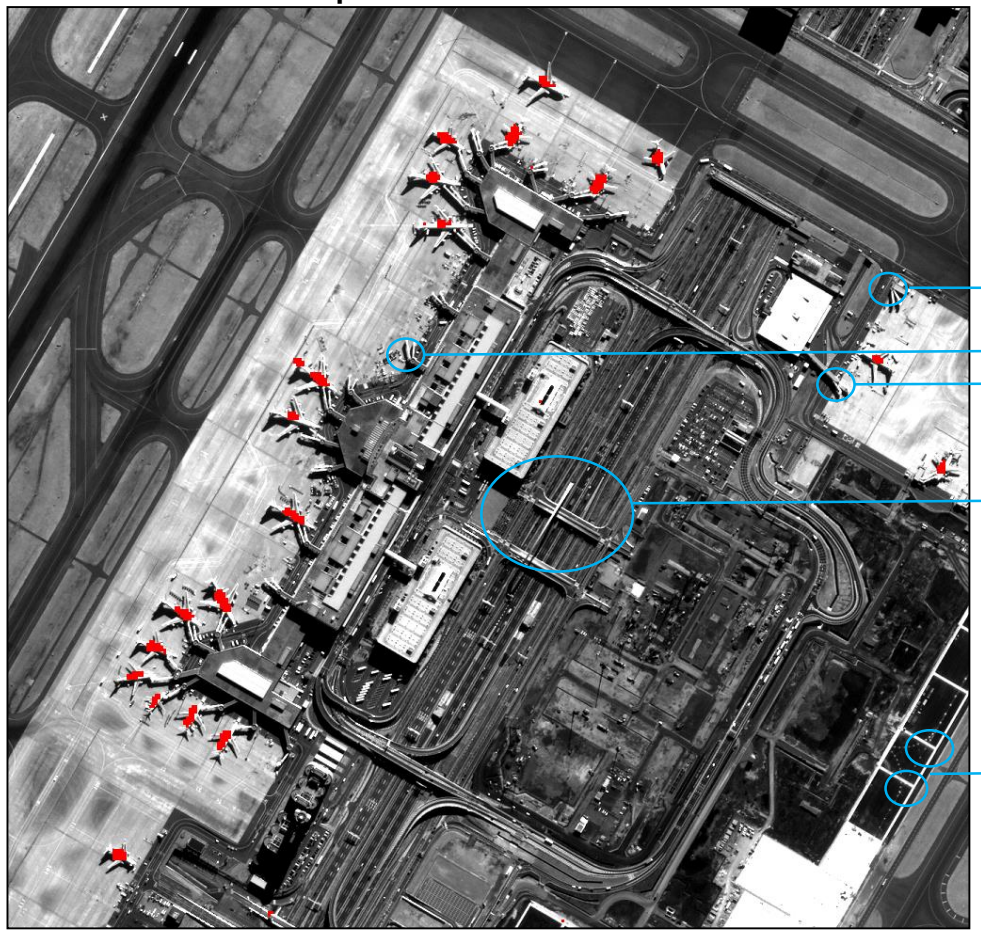

Padded & Stacked Inputs

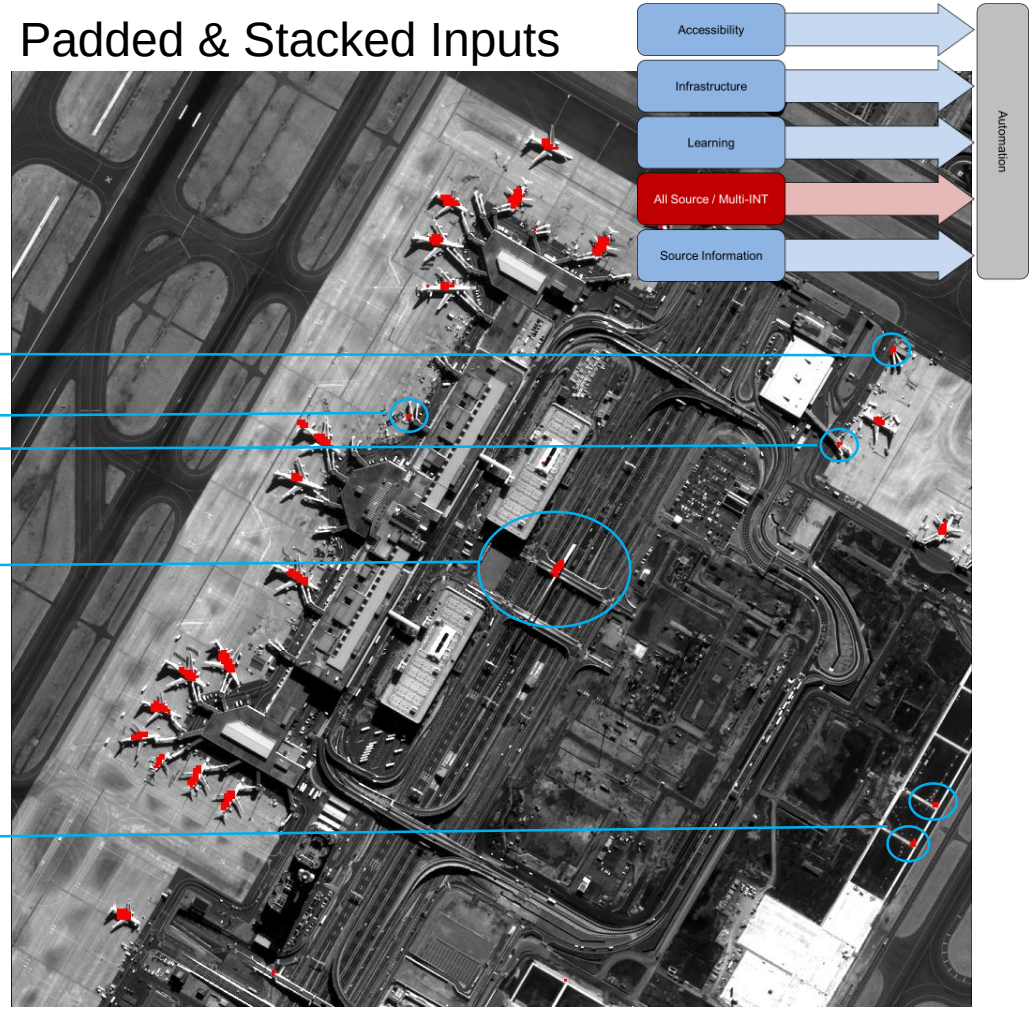

Persistent false positives overcome with use of derived data in multi-stream model with highly limited training data

Harris' work to scale deep learning for defense learning – transfer learning and domain adaptation

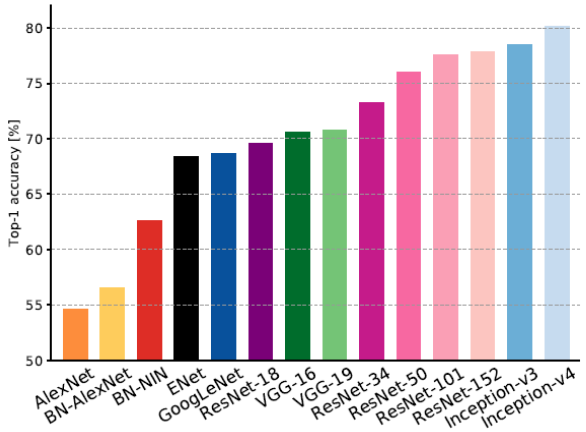

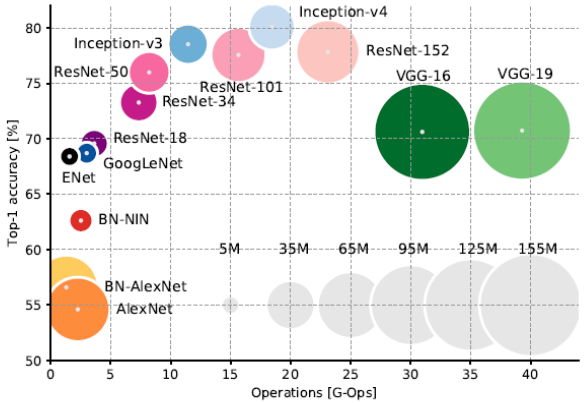

Source: arXiv:1605.07678, Canziani

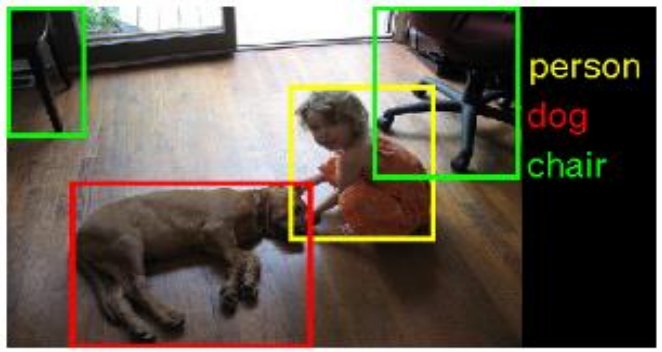

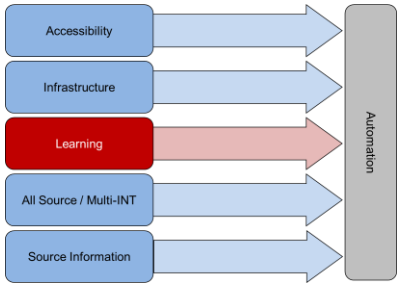

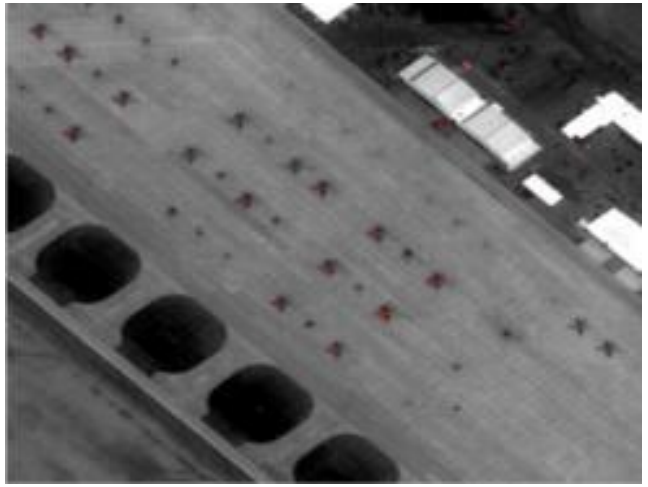

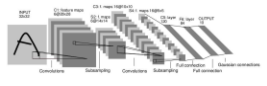

LeNet5
(1994 – 3 layers)

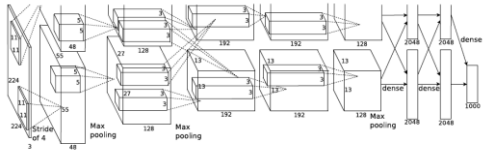

AlexNet
(2012 – 8 layers)

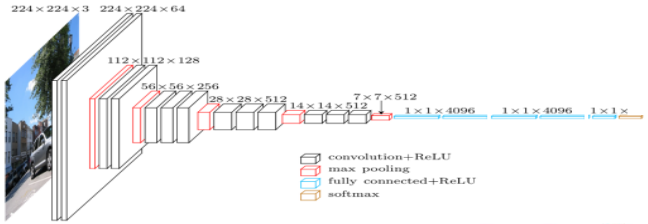

VGG16
(2014 – 16 layers)

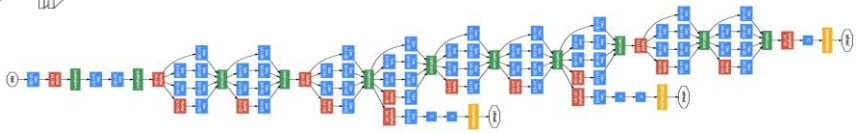

GoogLeNet
(2015 – 22 layers)

ResNet
(2015 – 152 layers)

Harris' work to scale deep learning for defense learning – improvements to conv-autoencoder for pretraining

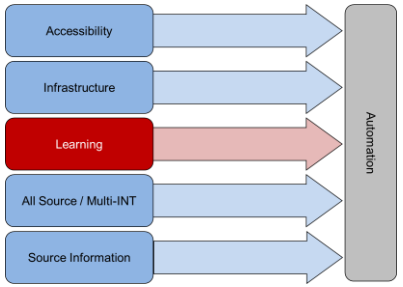

One of the few unsupervised methods for training CNNs. Can they be as discriminative as supervised training?

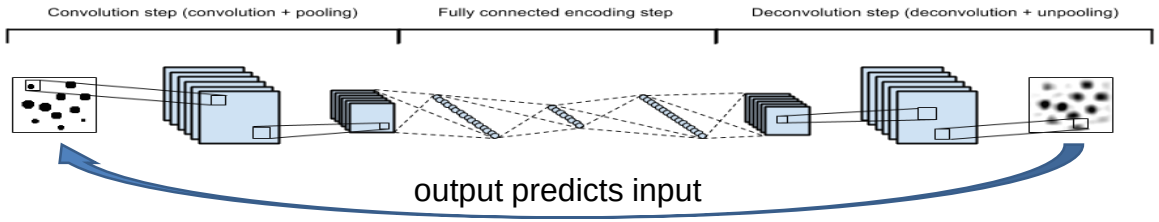

- When conv-autoencoder is applied to satellite imagery tend to learn simple convolutional features
- Far less effective for transfer learning then supervised classification CNNs
- With modifications, can adapt proven natural image network architectures to use conv-autoencoder for unsupervised pretraining with satellite imagery

Improvements in Low-level Conv Filters

VGG16 (ImageNet training: 1.2M samples, 1000 classes), Conv-0

Typical "Winner-Take-All" Conv-AE on Satellite RGB imagery – poor Conv-0 structure

Improved Harris Conv-AE on Satellite RGB imagery – better Conv-0 structure, closer to VGG16
Mods: Train all layers together, Enforce Sparsity, Selectively block gradient

Harris' work to scale deep learning for defense learning – AutoML

Blender – a tool for automatically building optimal neural network models from labeled data

- **Novices** can build a network without background knowledge
- **Experts** can explore more possibilities than by hand

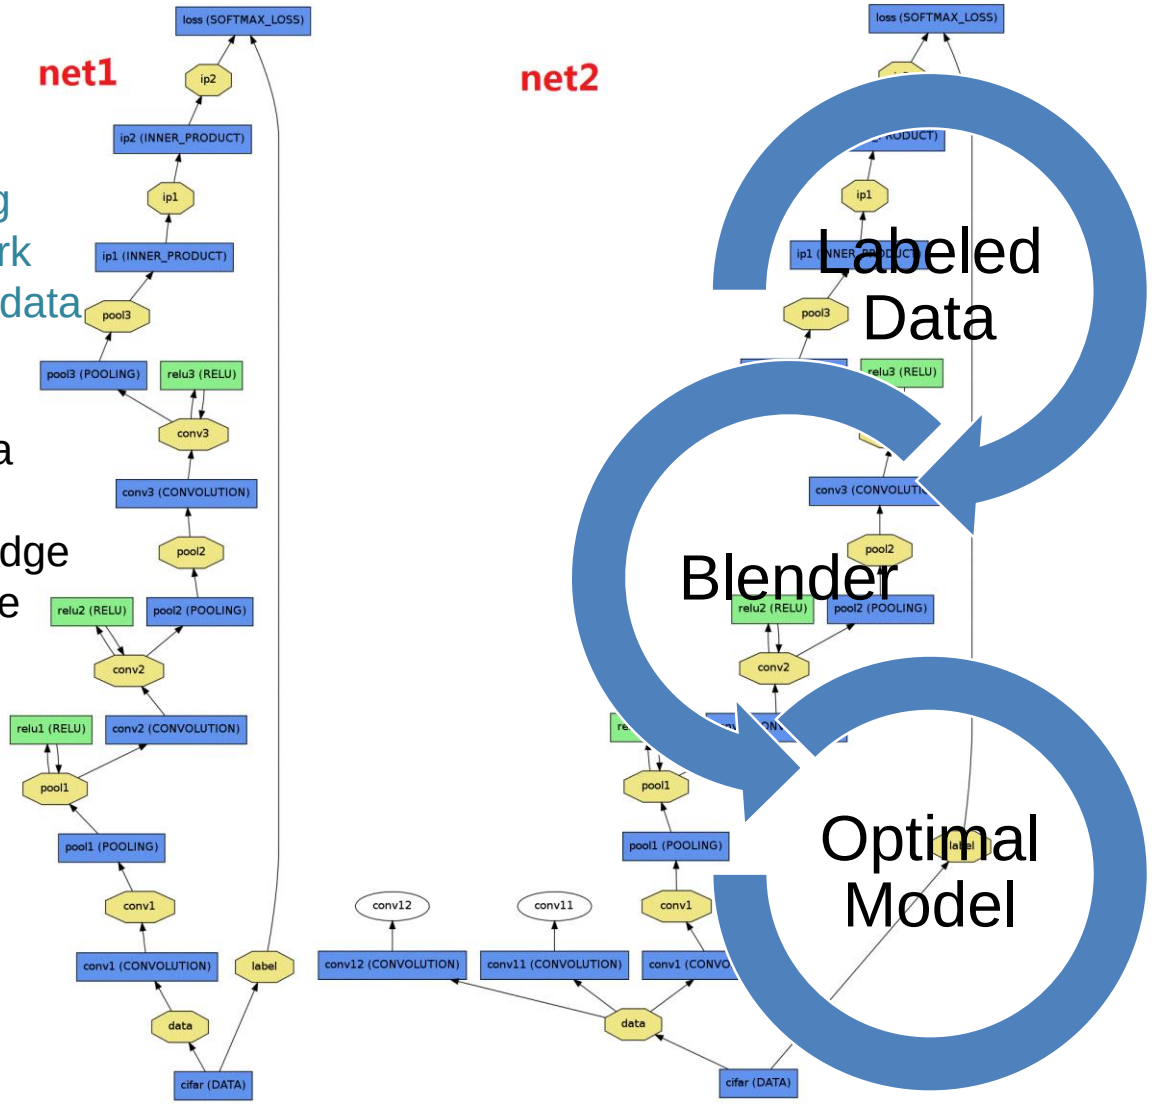

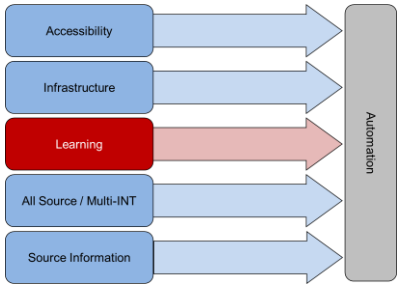

https://mathematica.stackexchange.com

Harris' work to scale deep learning for defense learning – AutoML

Best Human Model

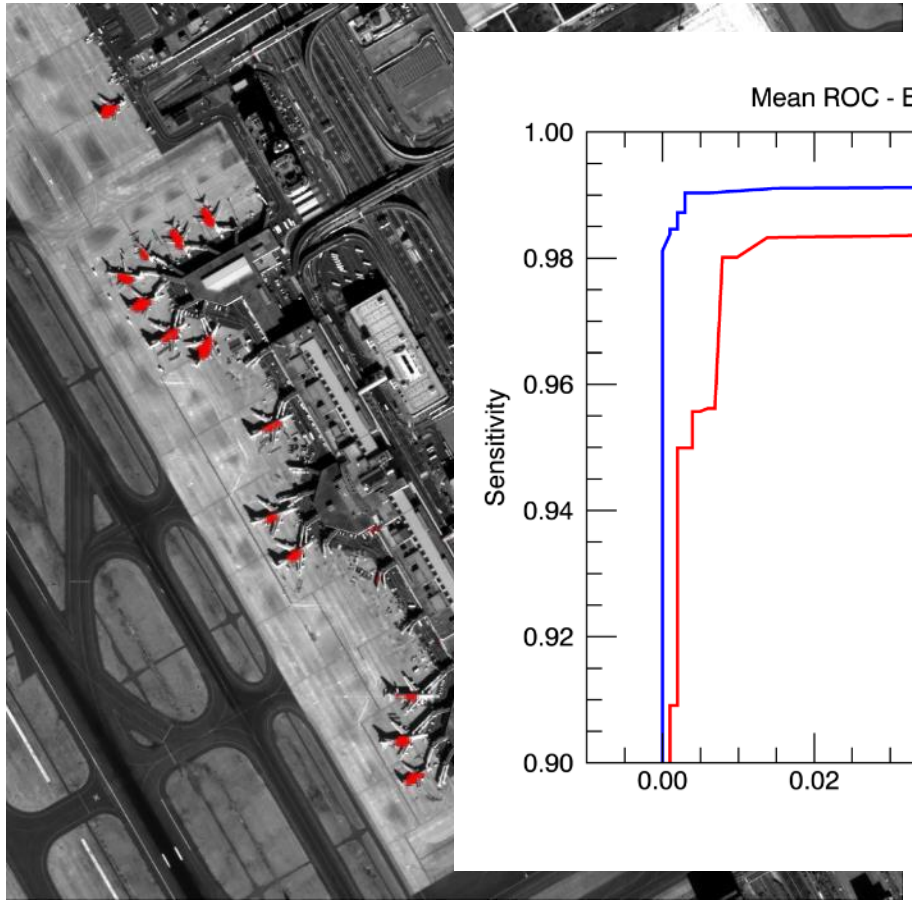

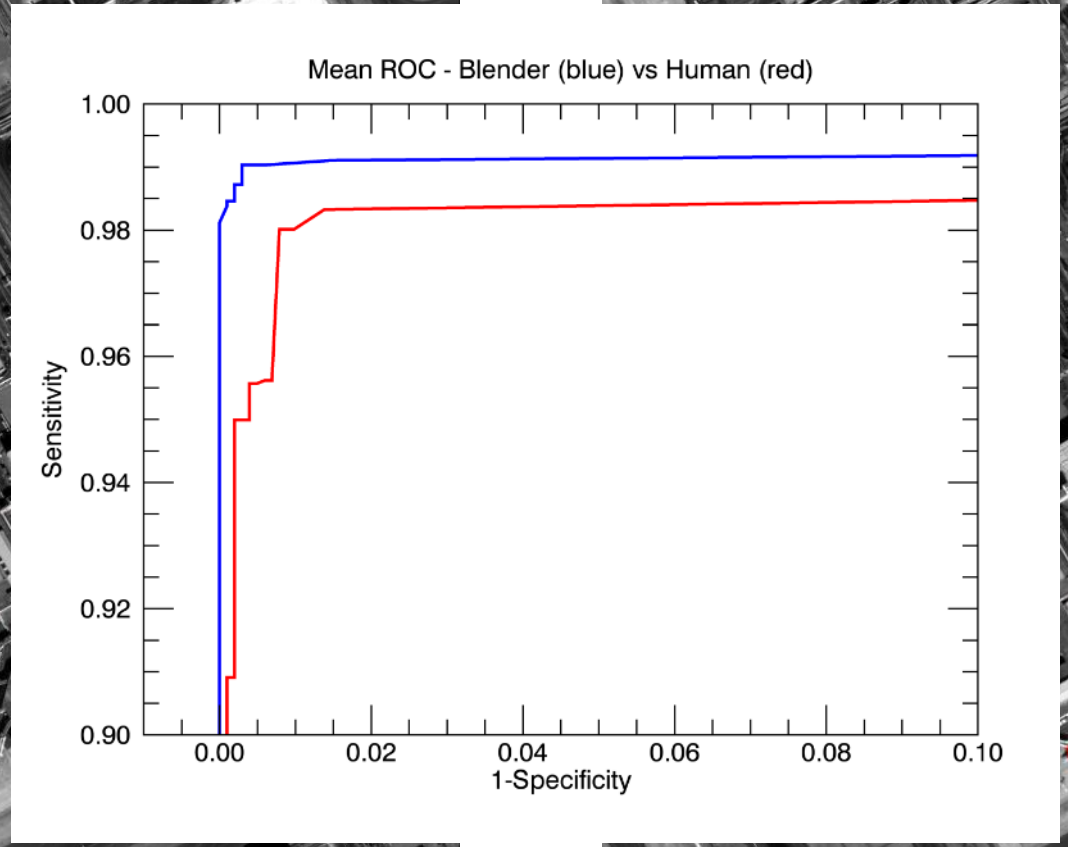

Best Blender Model

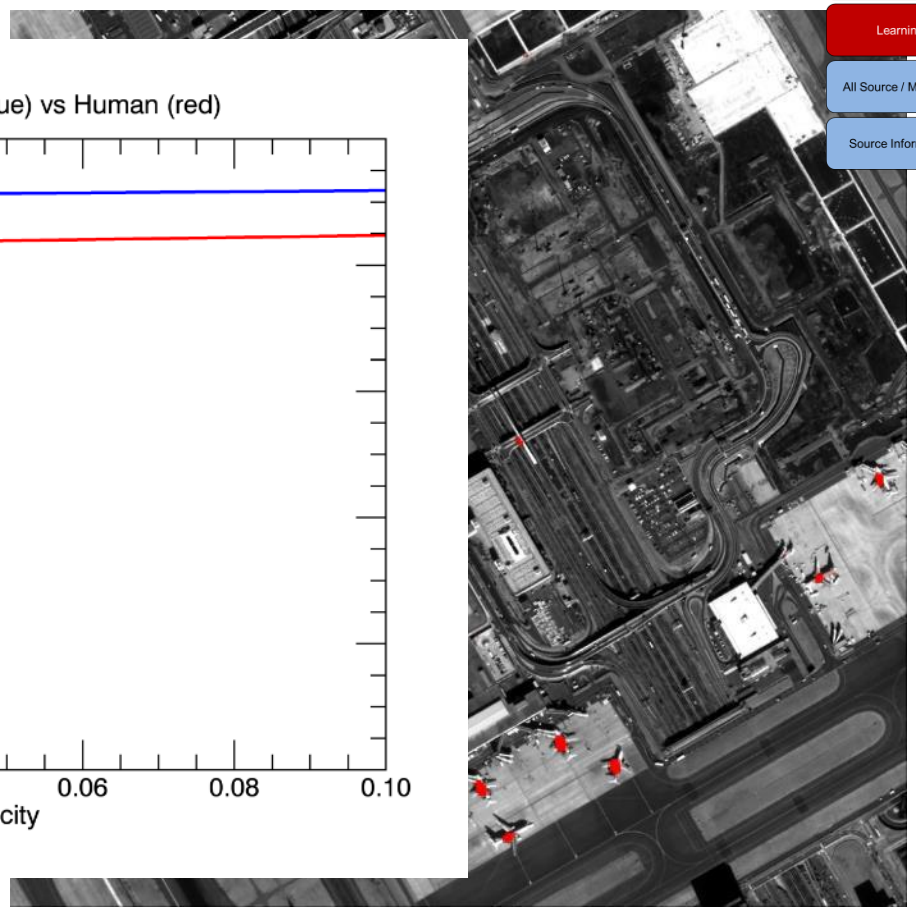

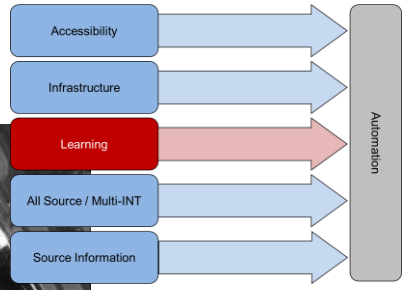

Harris' work to scale deep learning for defense infrastructure – exposing data repositories to deep learning

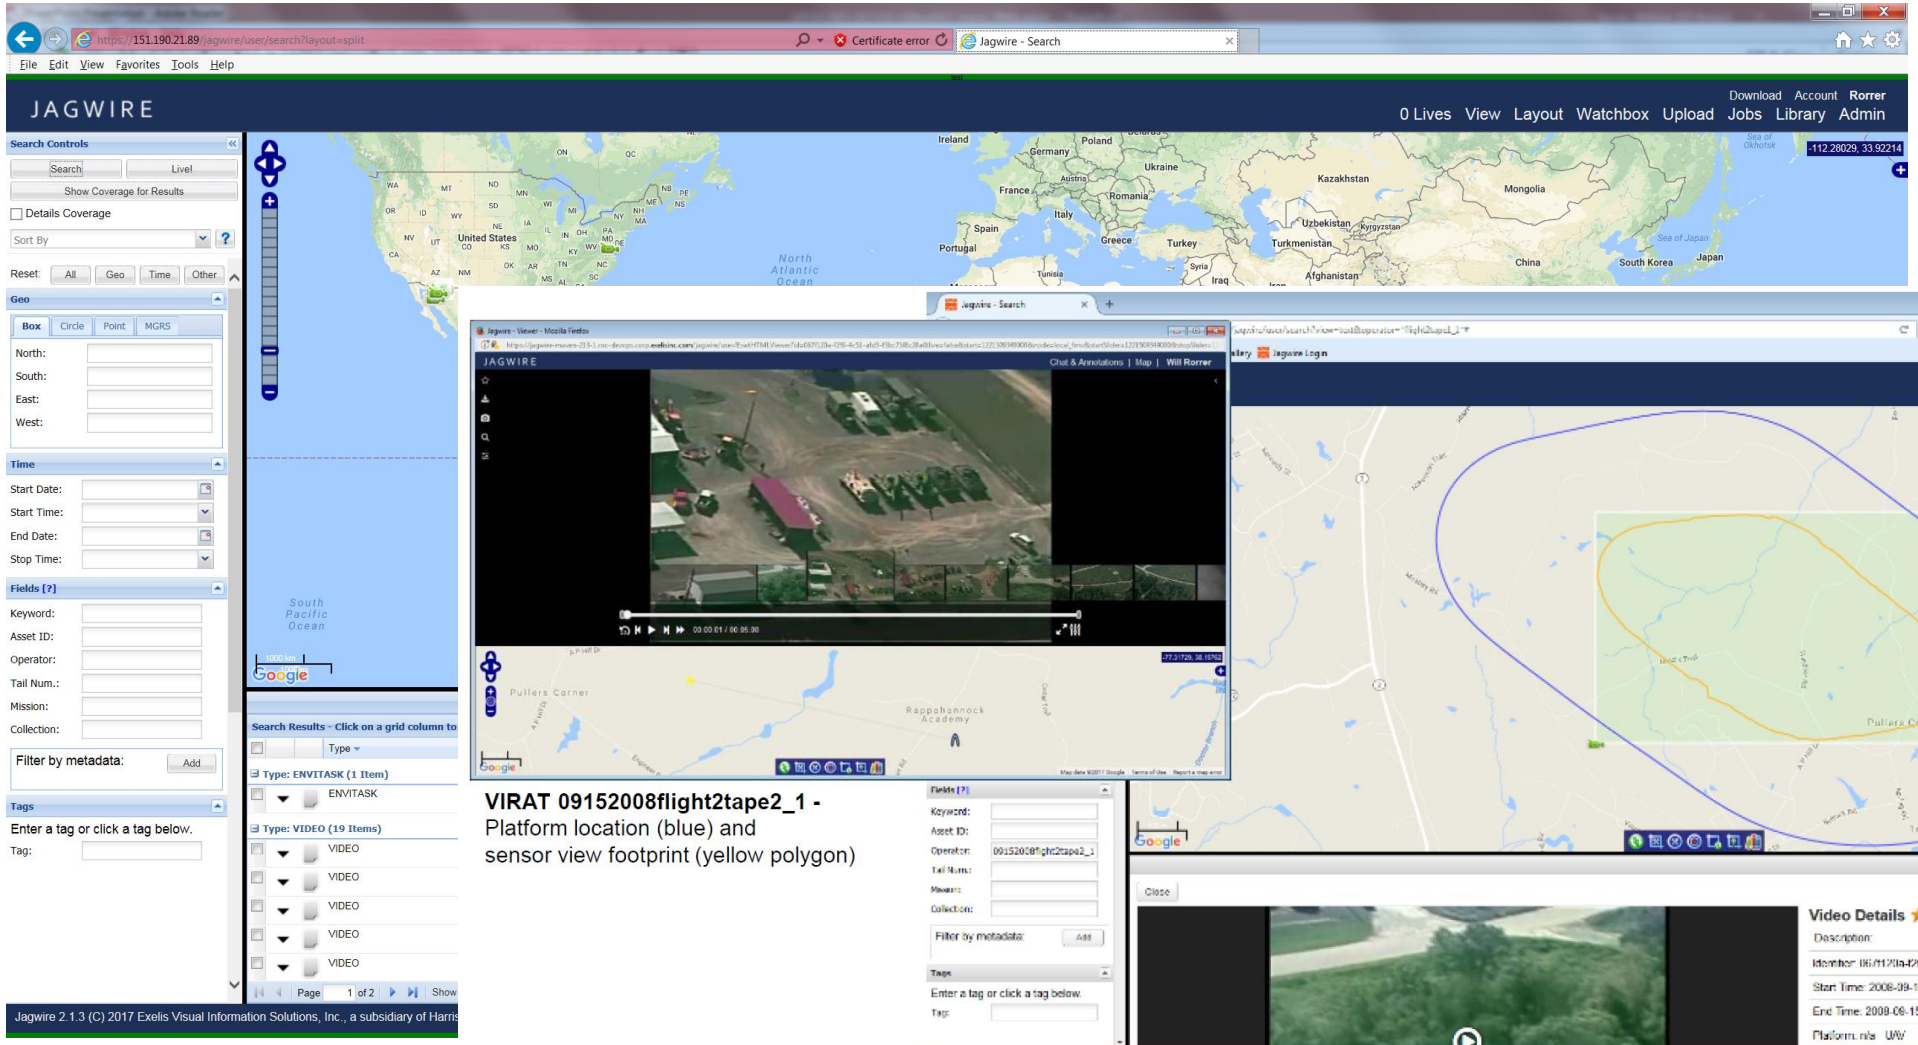

VIRAT 09152008flight2tape2_1 - Platform location (blue) and sensor view footprint (yellow polygon)

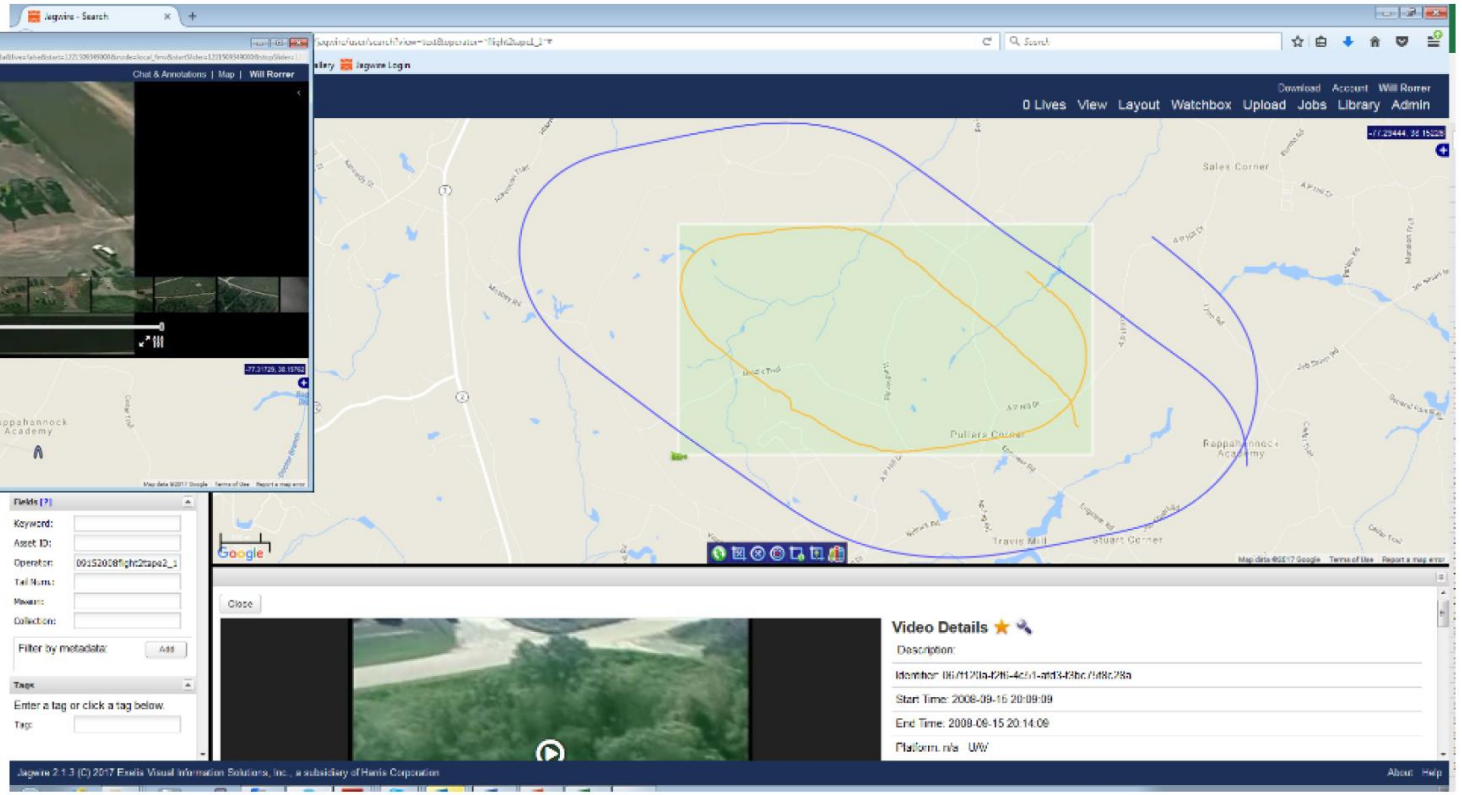

VIRAT 09152008flight2tape2_1 – Flight Path (Blue) and Sensor Coverage (Orange) Overview

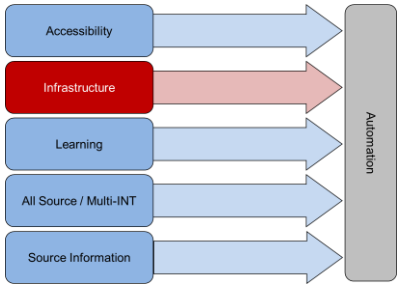

Harris' work to scale deep learning for defense infrastructure – incorporation of deep learning algorithms in workflows

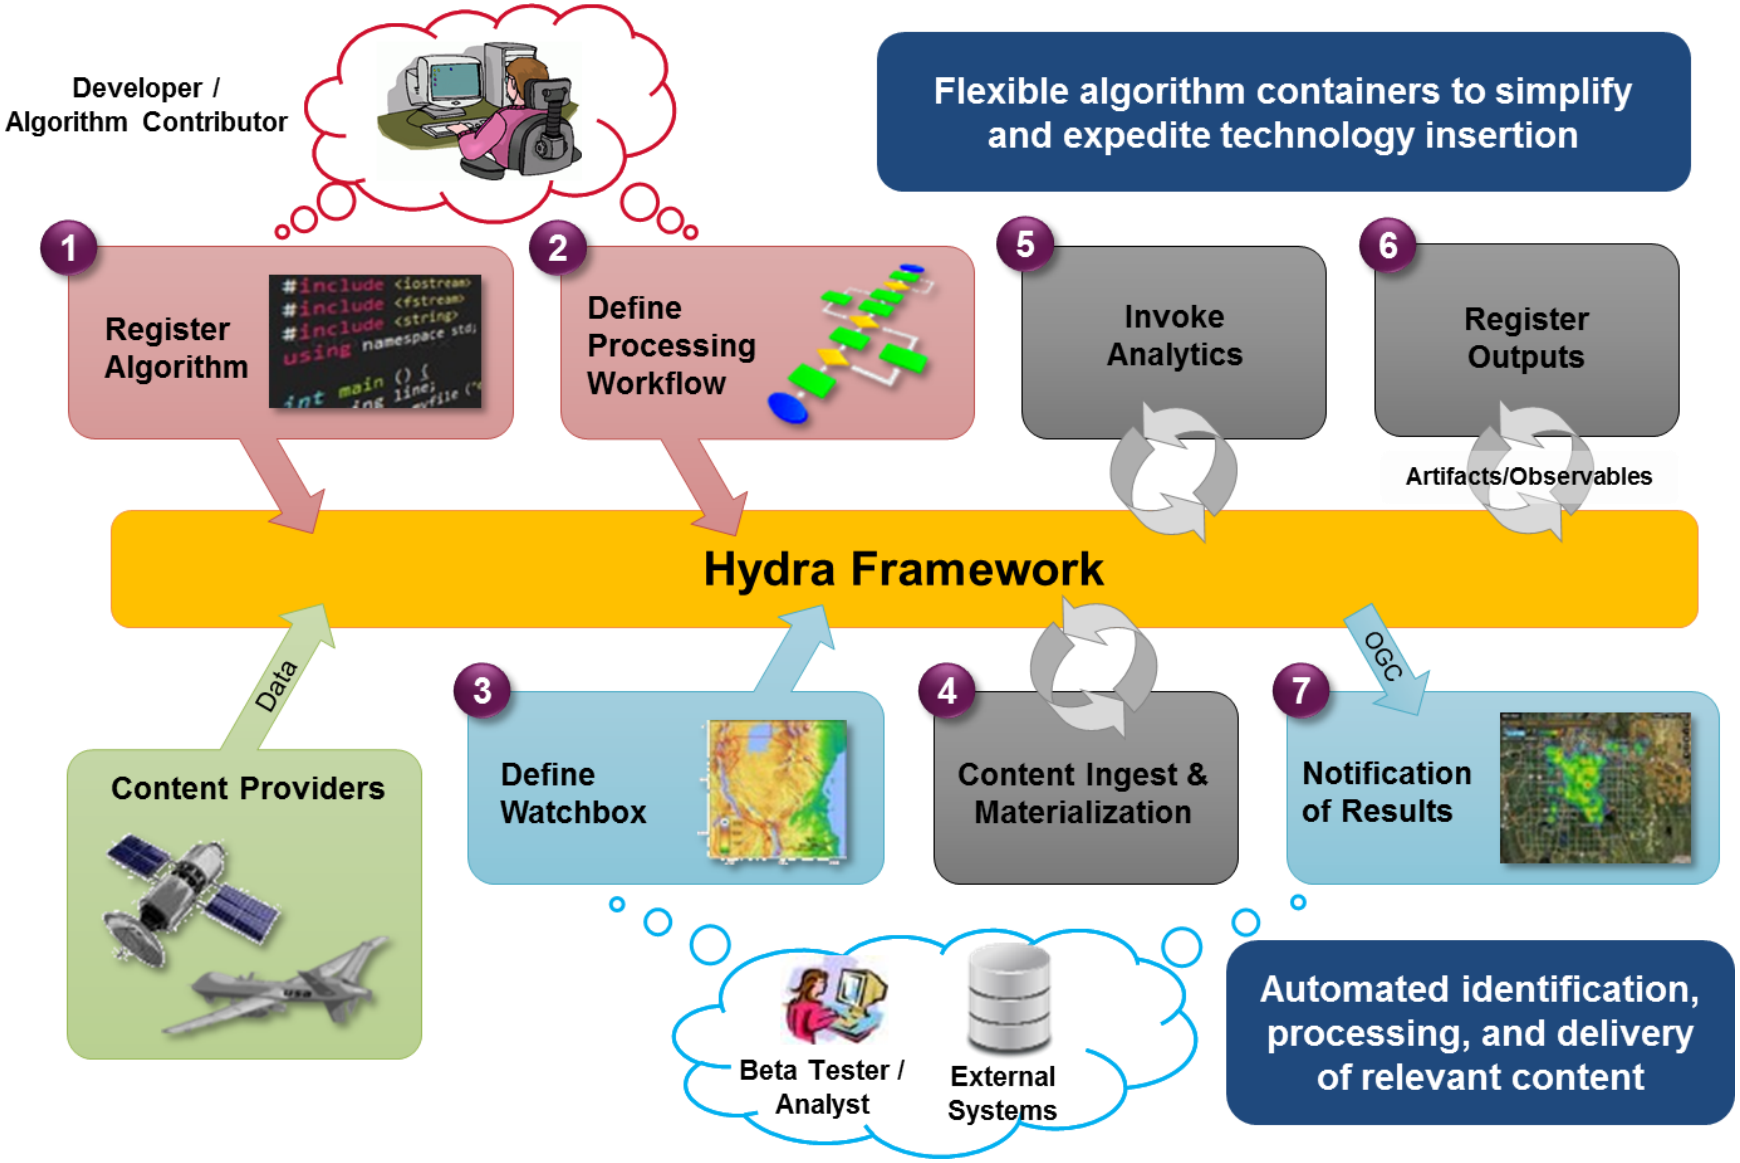

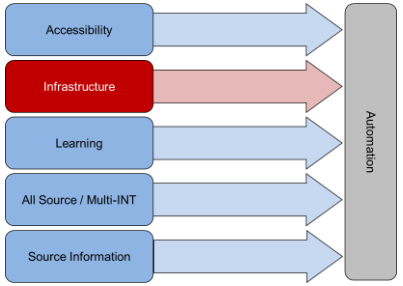

Harris' work to scale deep learning for defense infrastructure – incorporation of deep learning algorithms in workflows

- **Multiple algorithm "containers", or API's:**
  - Expedite development
  - Enable portability
  - Simplify integration
- **Algorithm Marketplace is API-neutral and provides:**
  - Streamlined integration and V&V
  - Web-based access to new products, algorithms, and services
  - Rapid development of analytic tools
  - Automatic and customized processing to address evolving customer needs
  - Ontology to drive "smart" algorithm execution

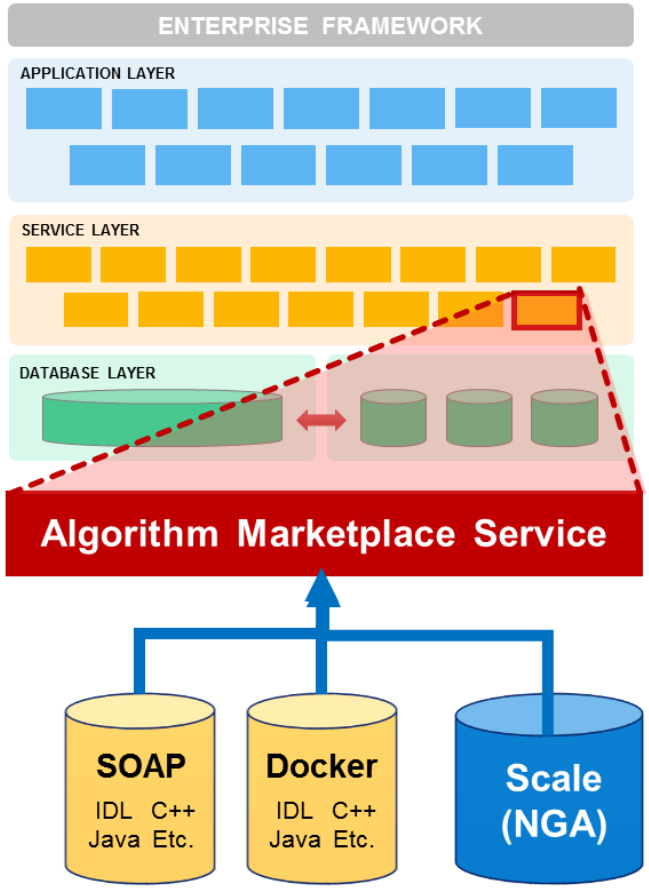

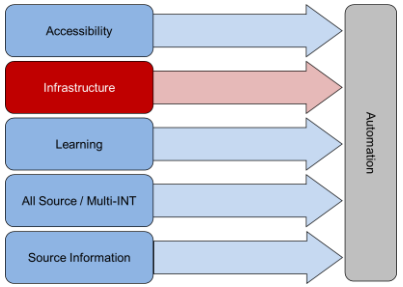

New analytics <u>rapidly inserted</u> via Algorithm Marketplace API's

Harris' work to scale deep learning for defense infrastructure – analyst assisted feedback loop

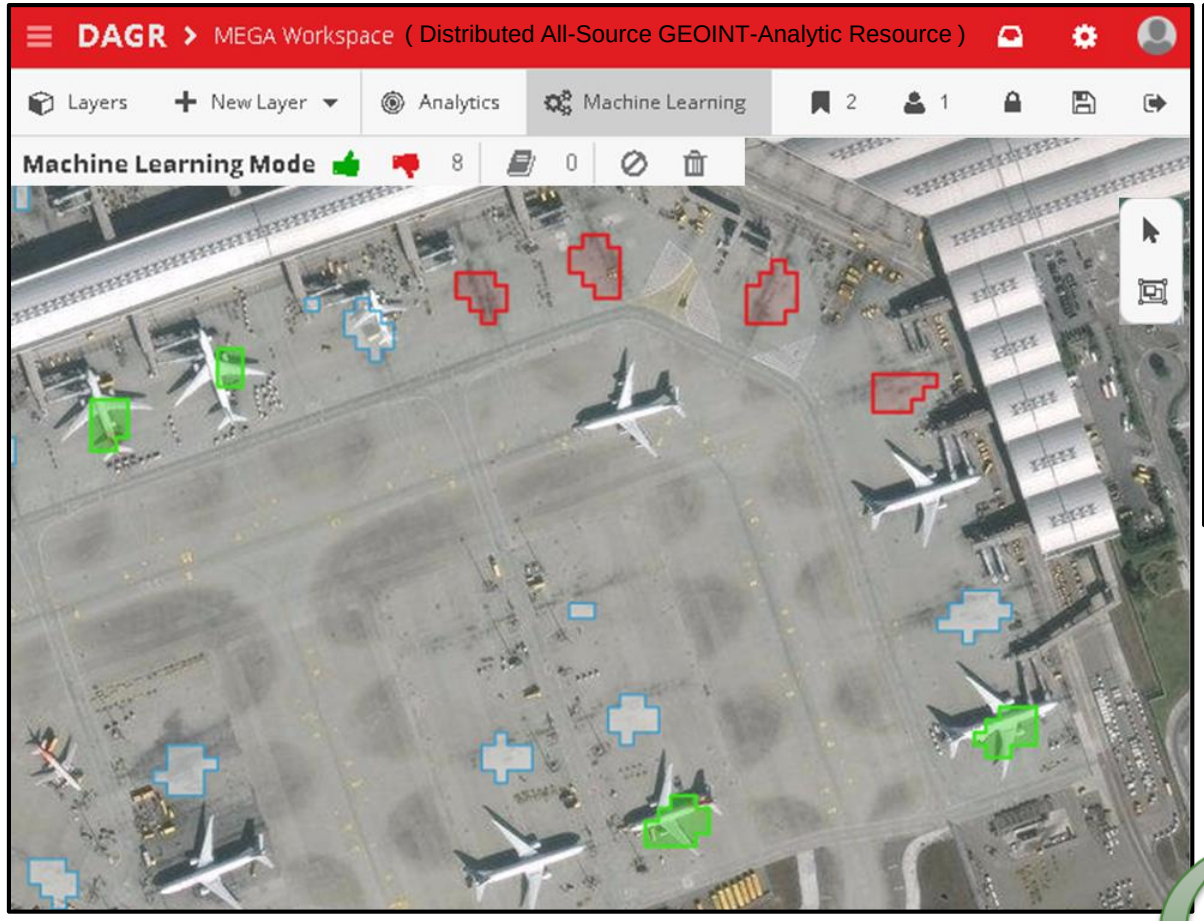

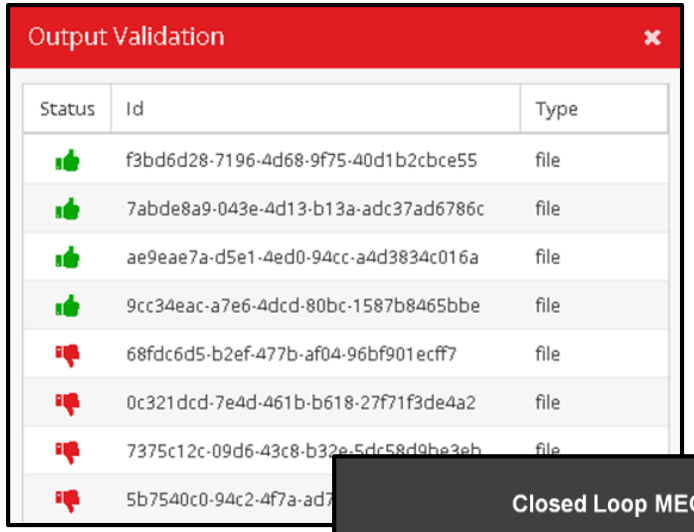

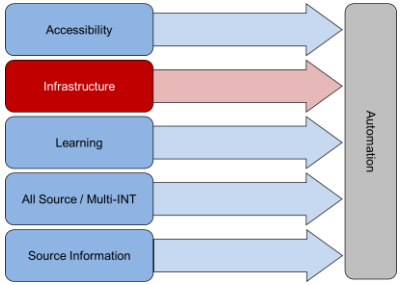

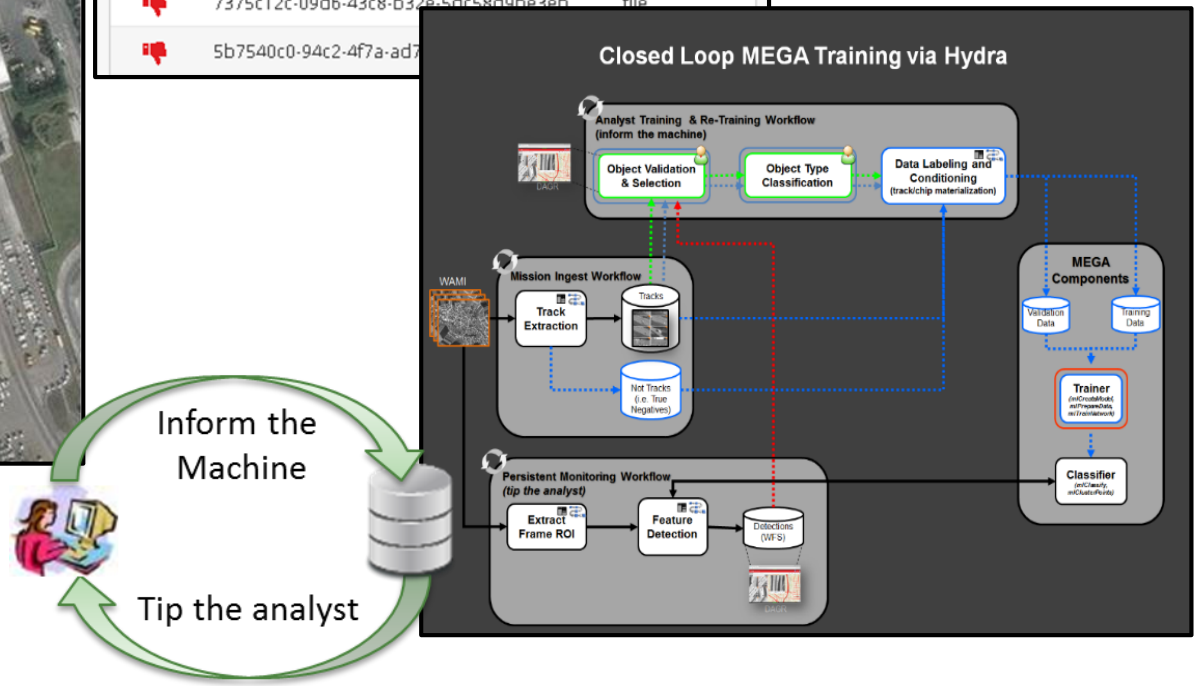

Harris' work to scale deep learning for defense accessibility – analyst-focused UI – flexible visualization

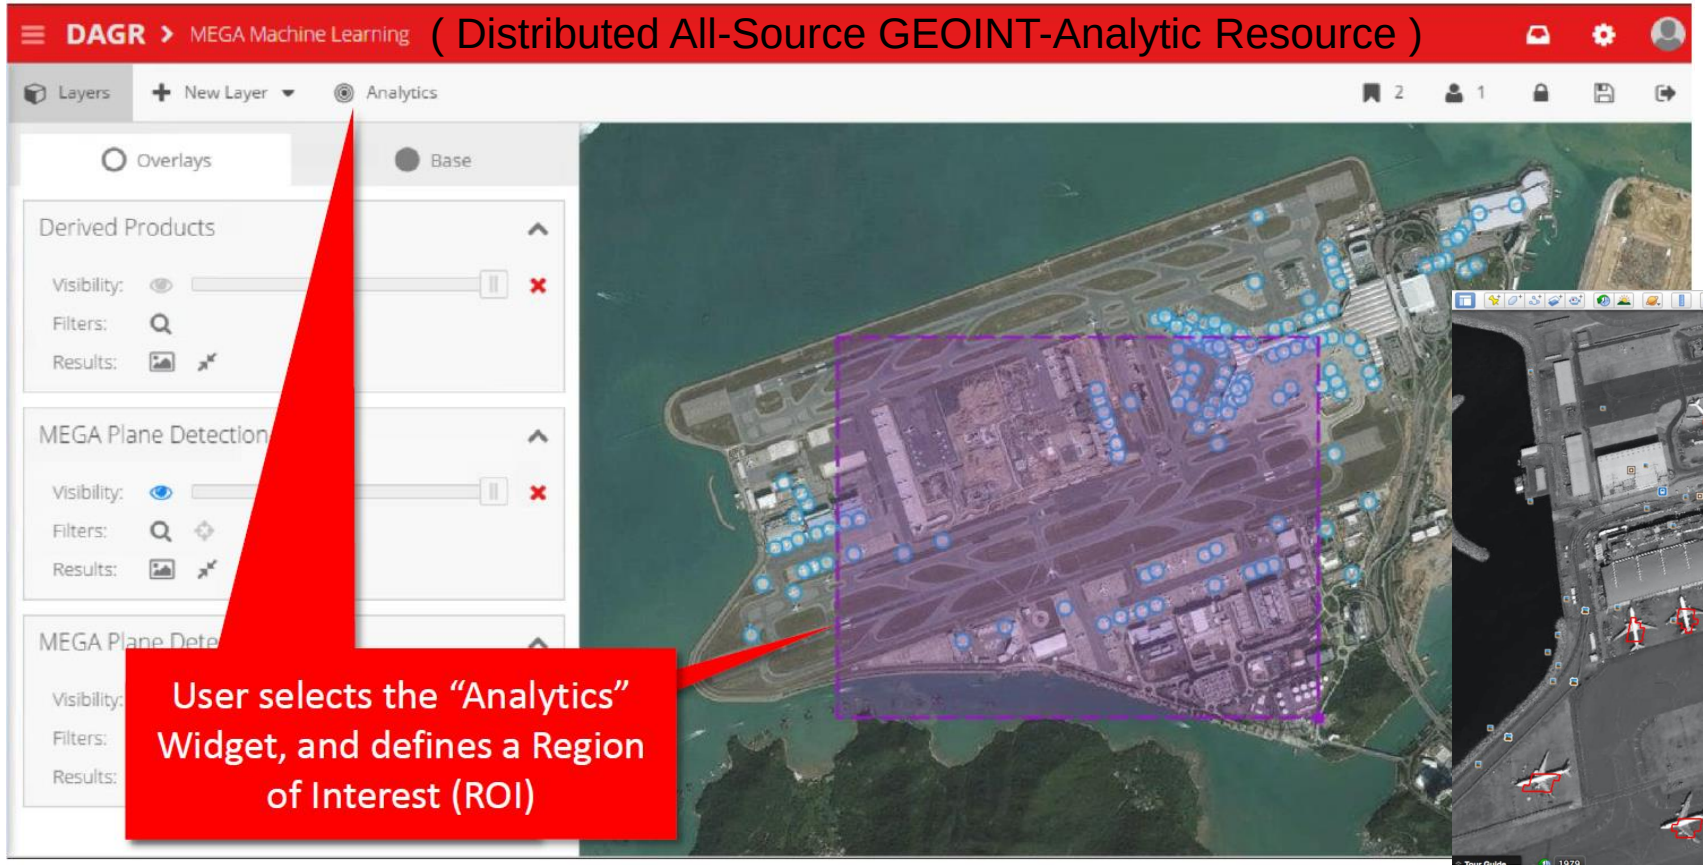

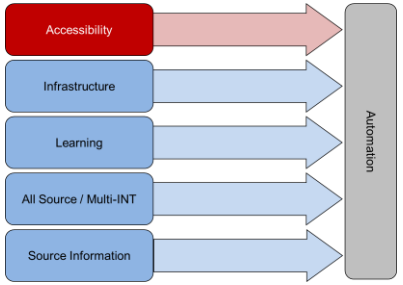

Map-based Analytics Recommendation available via Hydra-API for 3rd Party Client Integration

Harris' work to scale deep learning for defense accessibility – case study – airport traffic monitoring

Goal – Identify broad trends in air traffic

- Deep learning a critical building block
  - Object ID: find airplanes on tarmac
- Combined with other methods to make monitoring feasible
  - Too time consuming for analysts to look at every image
  - Need a way to compress information
- Aggregation of multiple detection algorithm results over time to identify patterns and anomalies in behaviour
  - Applicable to many similar monitoring problems

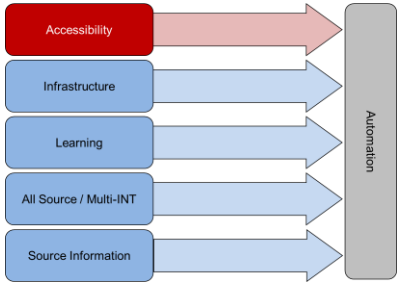

Test Case:

- Data: WorldView-3 Pan
- Airport: London Southend
- Multiple observations over several months

Object Detection

- LeNet-style 5-layer model converted for fully convolutional operation
- Detects all large airplanes

Harris' work to scale deep learning for defense accessibility – dashboard views

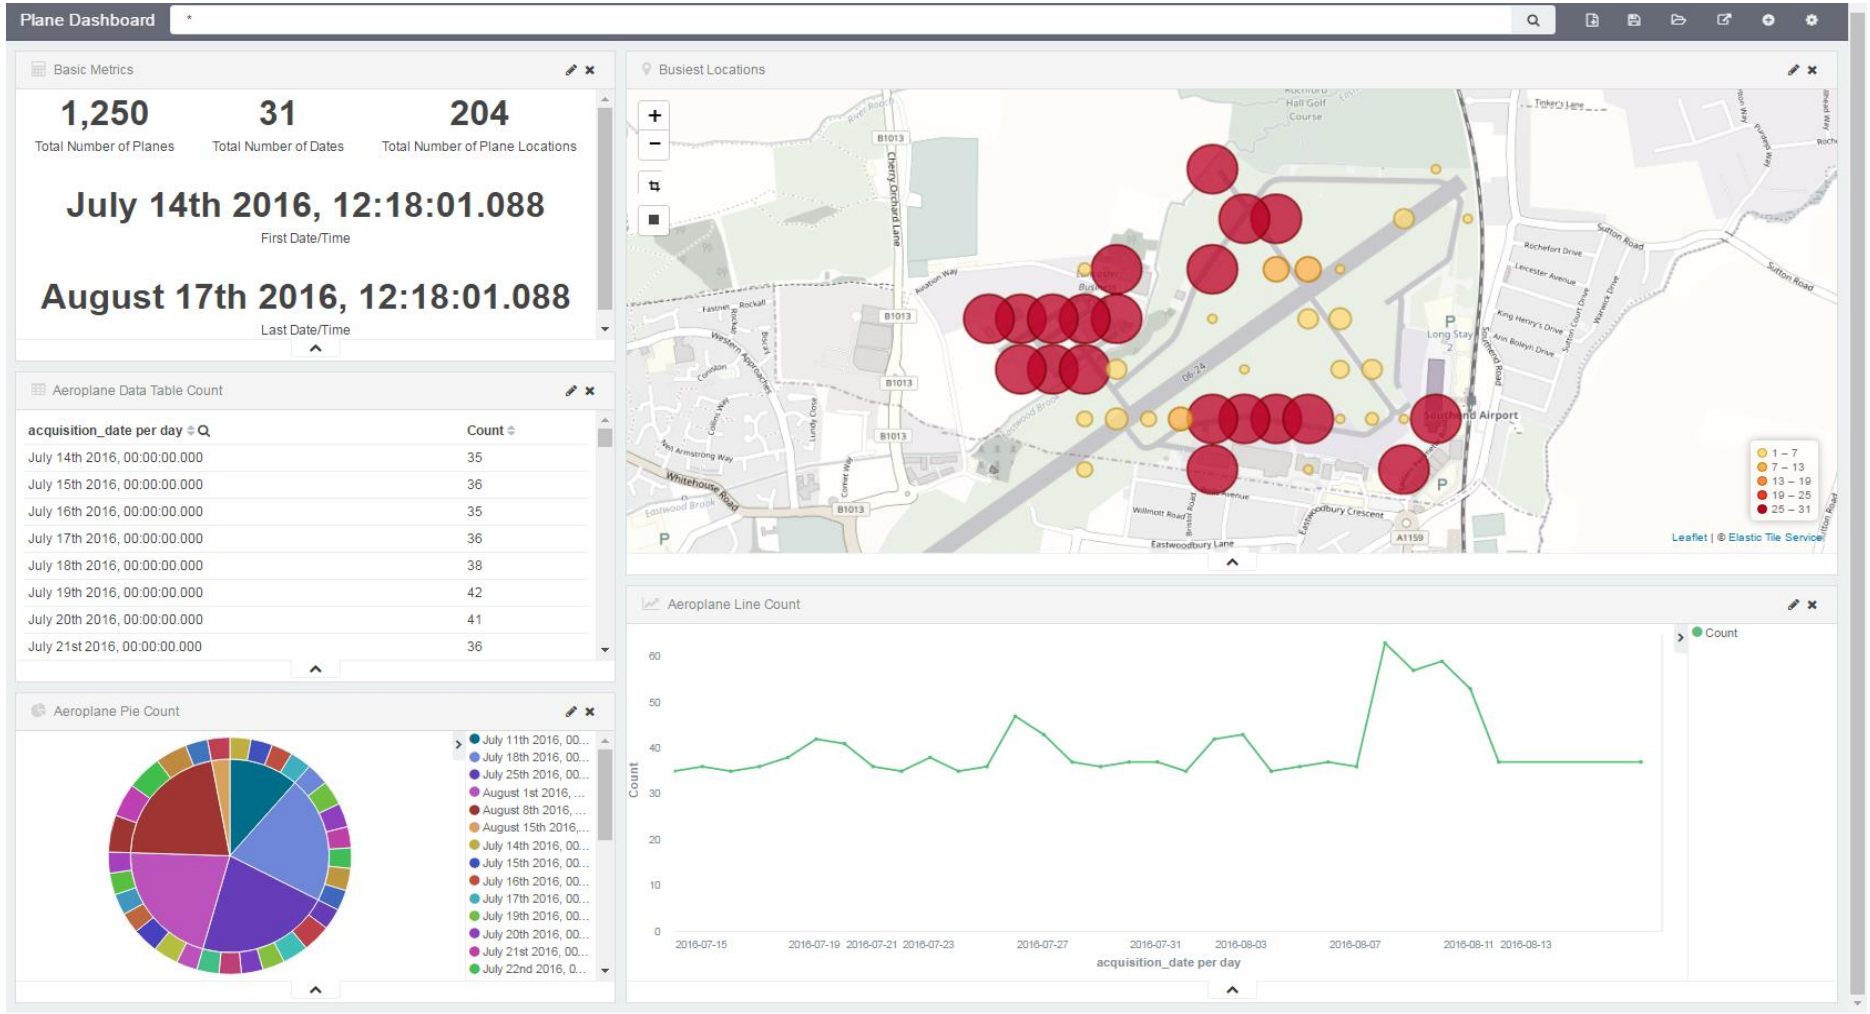

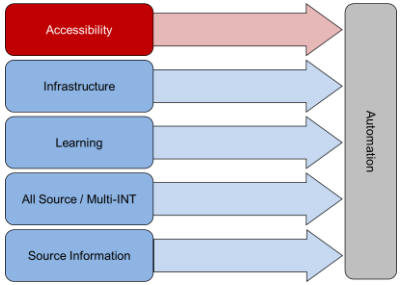

Harris' work to scale deep learning for defense automation – repeatable process

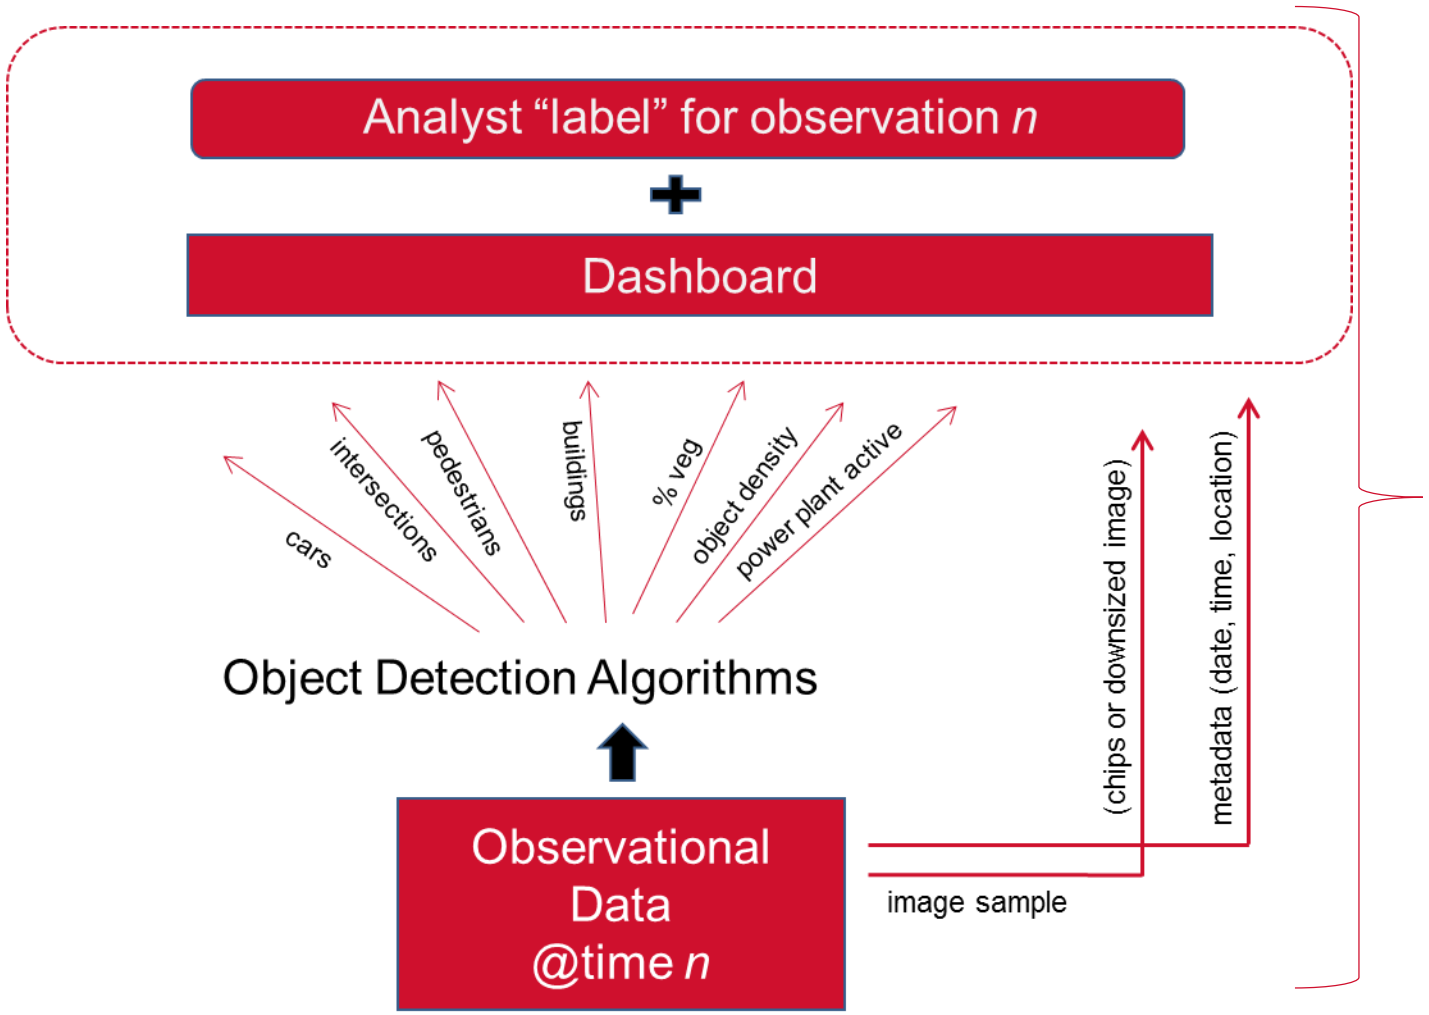

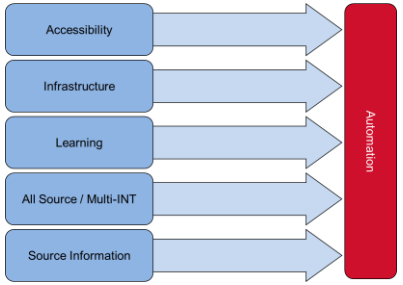

New training sample for decision support model is collected each time an analyst looks at dashboard and assigns a label, such as:

- Condition normal
- Emergency response
- Snowstorm
- Flood
- Power outage
- Traffic jam
- Suspicious activity

Harris's work to scale deep learning for defense automation – higher order sense making

In the not too distant future…

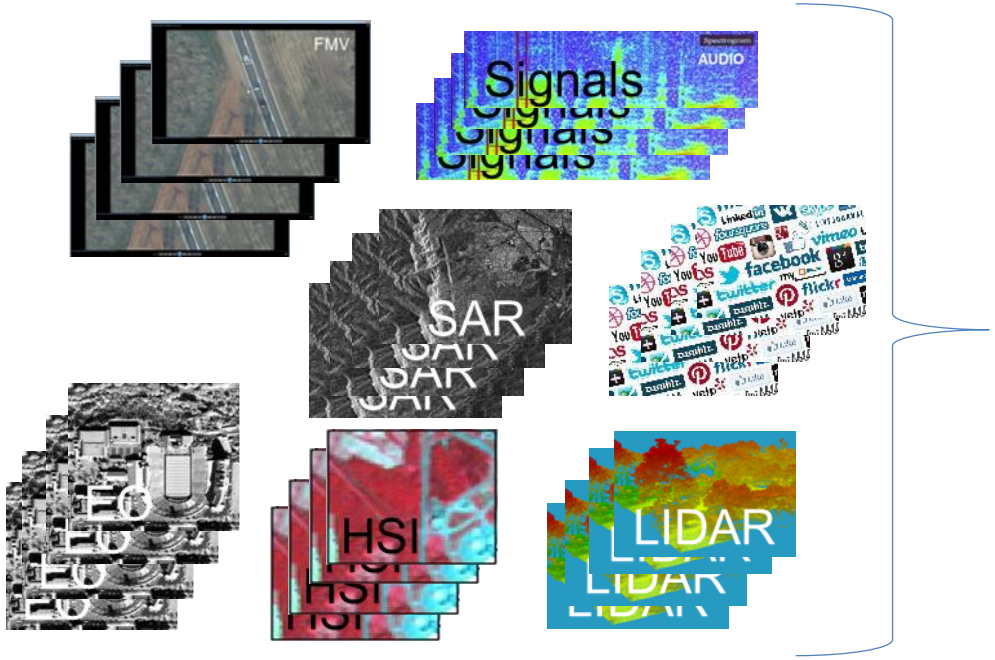

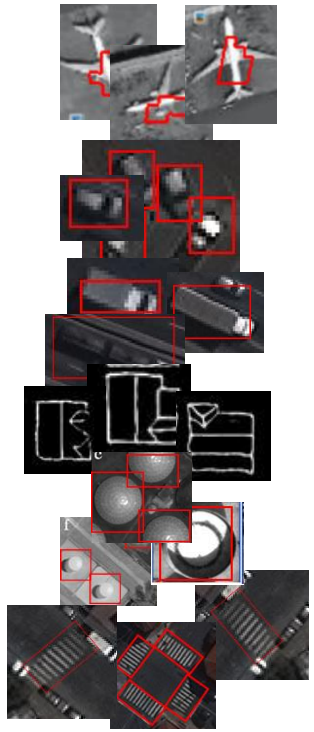

Live Observations Tally:

Planes Detected:	03
Cars Detected:	22
Trucks Detected:	10
Ships Detected:	02
Storage Tanks Detected:	08
Buildings Detected:	05
Crosswalks Detected:	12
Intersections Detected:	15
Etc…..	

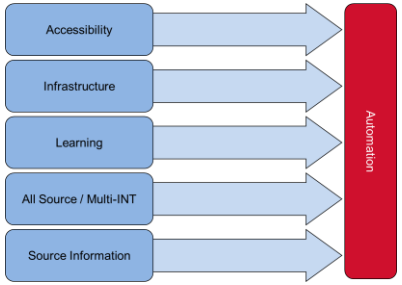

We will go from drowning in data…

…to drowning in detections/features/observations.

Deep learning/machine learning/artificial intelligence attention will shift from detecting objects and features to detecting actions, events, patterns of life.

Conclusions/takeaway

Deep learning removed the bottlenecks of having to hand construct algorithms for target/feature/event/activity recognition

However, deep learning has its own bottlenecks:

- Label data availability
- Network architecture
- Algorithm brittleness / training data variation distribution
- Ability to quickly adapt to new domain and/or target

Eliminating / minimizing those bottlenecks at scale is possible

Mass production of deep learning based GEOINT algorithms is within reach!

Questions?

Look for upcoming blog posts about machine learning at Harris' Blog website:
http://www.harrisgeospatial.com/Company/PressRoom/Blogs.aspx

Will Rorrer, PMP
Machine Learning
Business Development
william.rorrer@harris.com

Ron Kneusel, PHD
Senior Data Scientist
ronald.kneusel@harris.com

Backup Slides

Breaking it down

Infrastructure

- Security
- Hardware Optimization
- Archival
- Versioning
- APIs
- Componentization
- Open Source

Accessibility

- Analyst retraining
- Watchbox

All Source / Multi-INT

- Multistream
- Data support
- Data access
- Structured Data
- Unstructured data

Learning

- CNNs, RNNs, LSTMs…
- Engine support (Café, Theano, Tensorflow…)
- AutoML

Training

- Transfer Learning
- Crowdsourcing
- Modeling
- Analyst feedback

Automation

Cognitive Ecosystem

Hydra applications
Collaborative content management and analytical processing

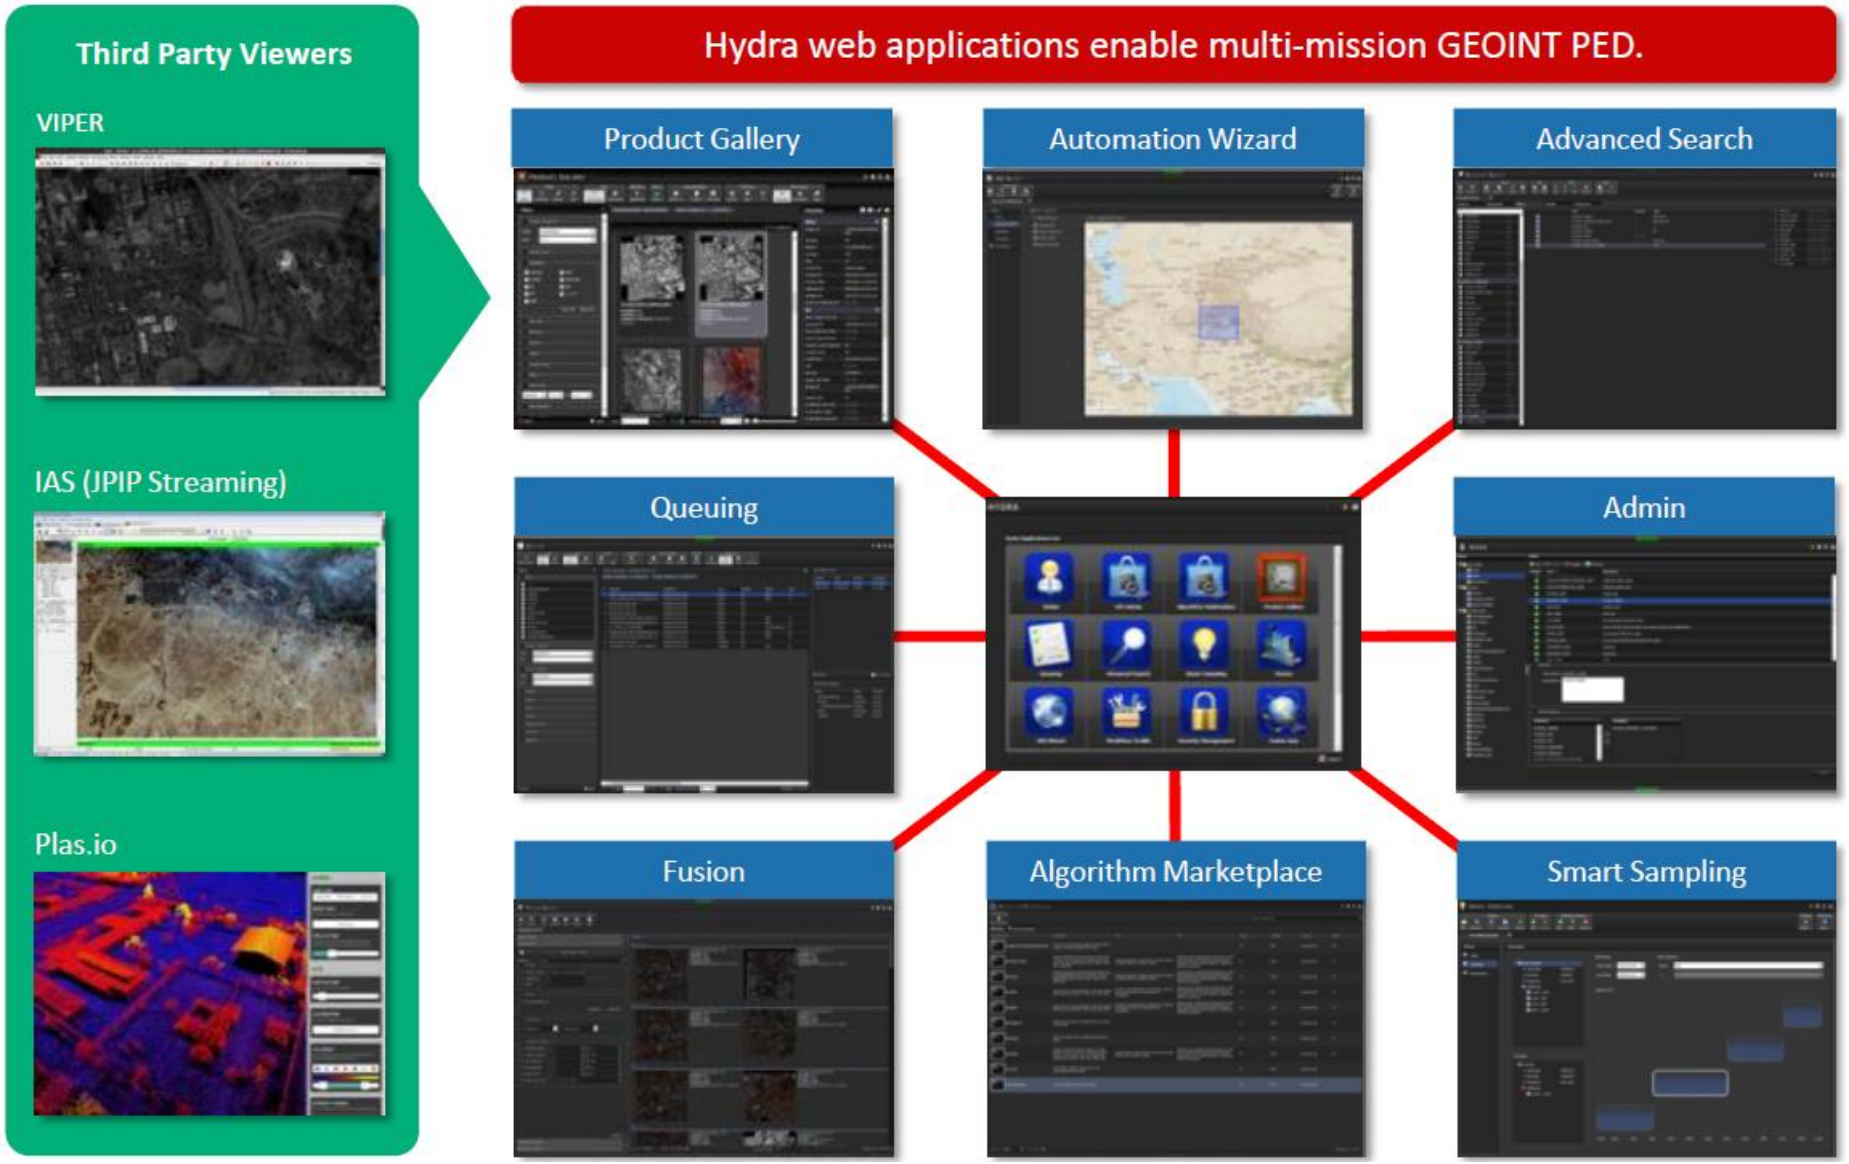

Rapid insertion of analytics and self-service processing

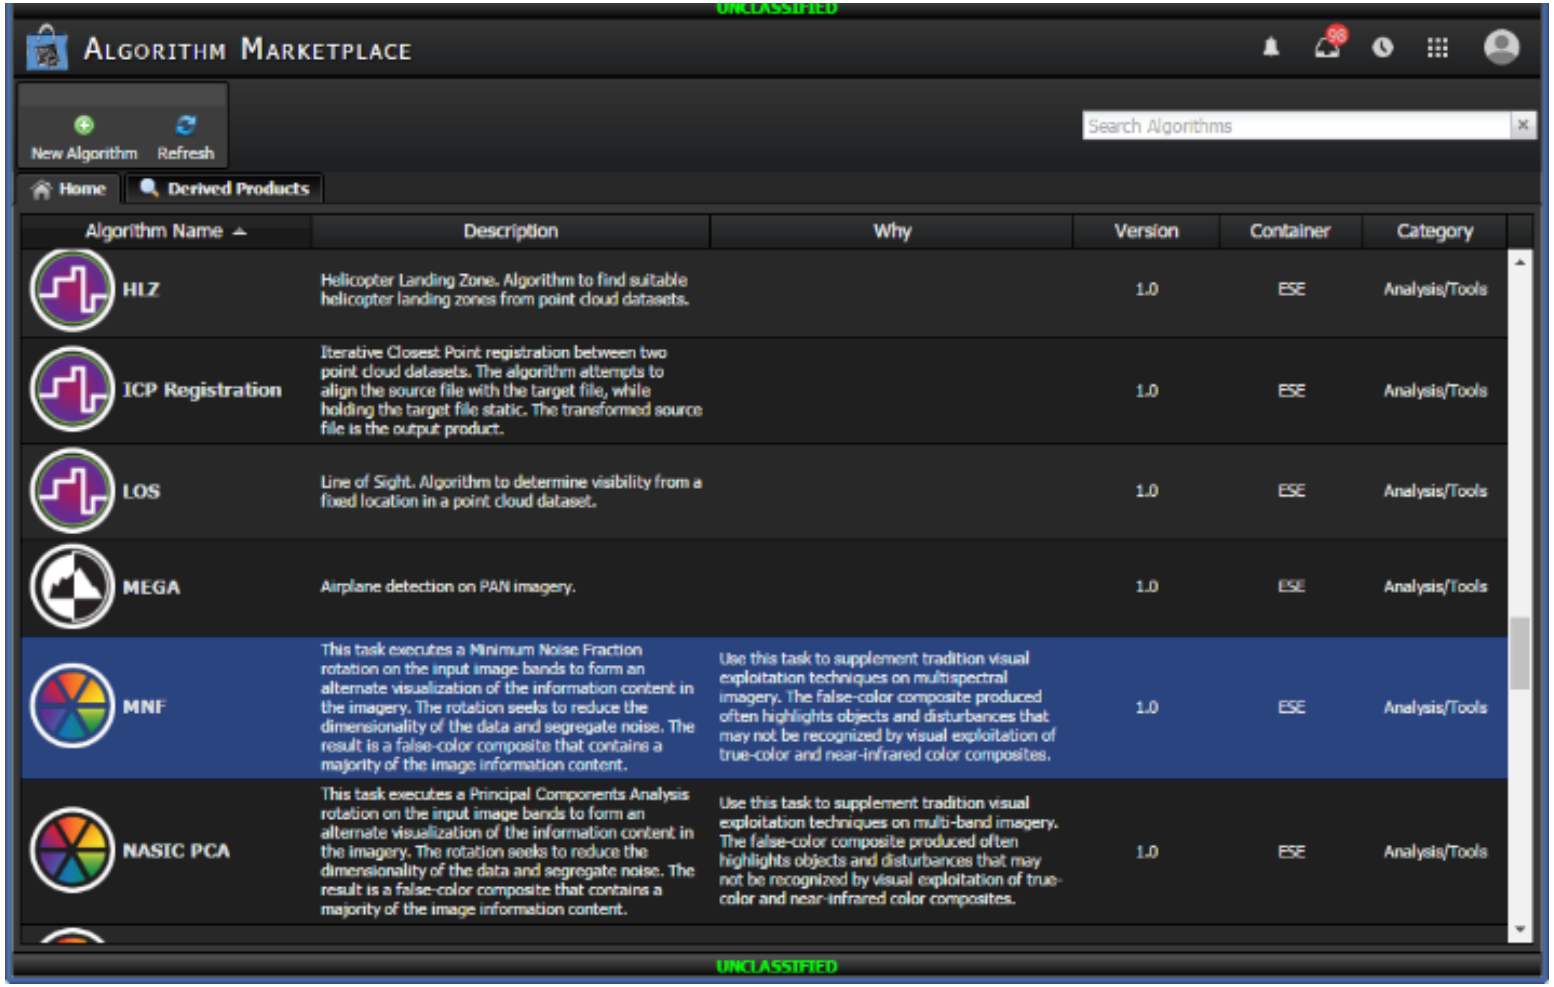

Algorithm Name	Description	Why	Version	Container	Category
HLZ	Helicopter Landing Zone. Algorithm to find suitable helicopter landing zones from point cloud datasets.		1.0	ESE	Analysis/Tools
ICP Registration	Iterative Closest Point registration between two point cloud datasets. The algorithm attempts to align the source file with the target file, while holding the target file static. The transformed source file is the output product.		1.0	ESE	Analysis/Tools
LOS	Line of Sight. Algorithm to determine visibility from a fixed location in a point cloud dataset.		1.0	ESE	Analysis/Tools
MEGA	Airplane detection on PAN imagery.		1.0	ESE	Analysis/Tools
MNF	This task executes a Minimum Noise Fraction rotation on the input image bands to form an alternate visualization of the information content in the imagery. The rotation seeks to reduce the dimensionality of the data and segregate noise. The result is a false-color composite that contains a majority of the image information content.	Use this task to supplement tradition visual exploitation techniques on multispectral imagery. The false-color composite produced often highlights objects and disturbances that may not be recognized by visual exploitation of true-color and near-infrared color composites.	1.0	ESE	Analysis/Tools
NASIC PCA	This task executes a Principal Components Analysis rotation on the input image bands to form an alternate visualization of the information content in the imagery. The rotation seeks to reduce the dimensionality of the data and segregate noise. The result is a false-color composite that contains a majority of the image information content.	Use this task to supplement tradition visual exploitation techniques on multi-band imagery. The false-color composite produced often highlights objects and disturbances that may not be recognized by visual exploitation of true-color and near-infrared color composites.	1.0	ESE	Analysis/Tools

Multiple methods to register analytics enable rapid integration of new capabilities and drive minimal system integrator (SI) support

Analytics governance and collaboration

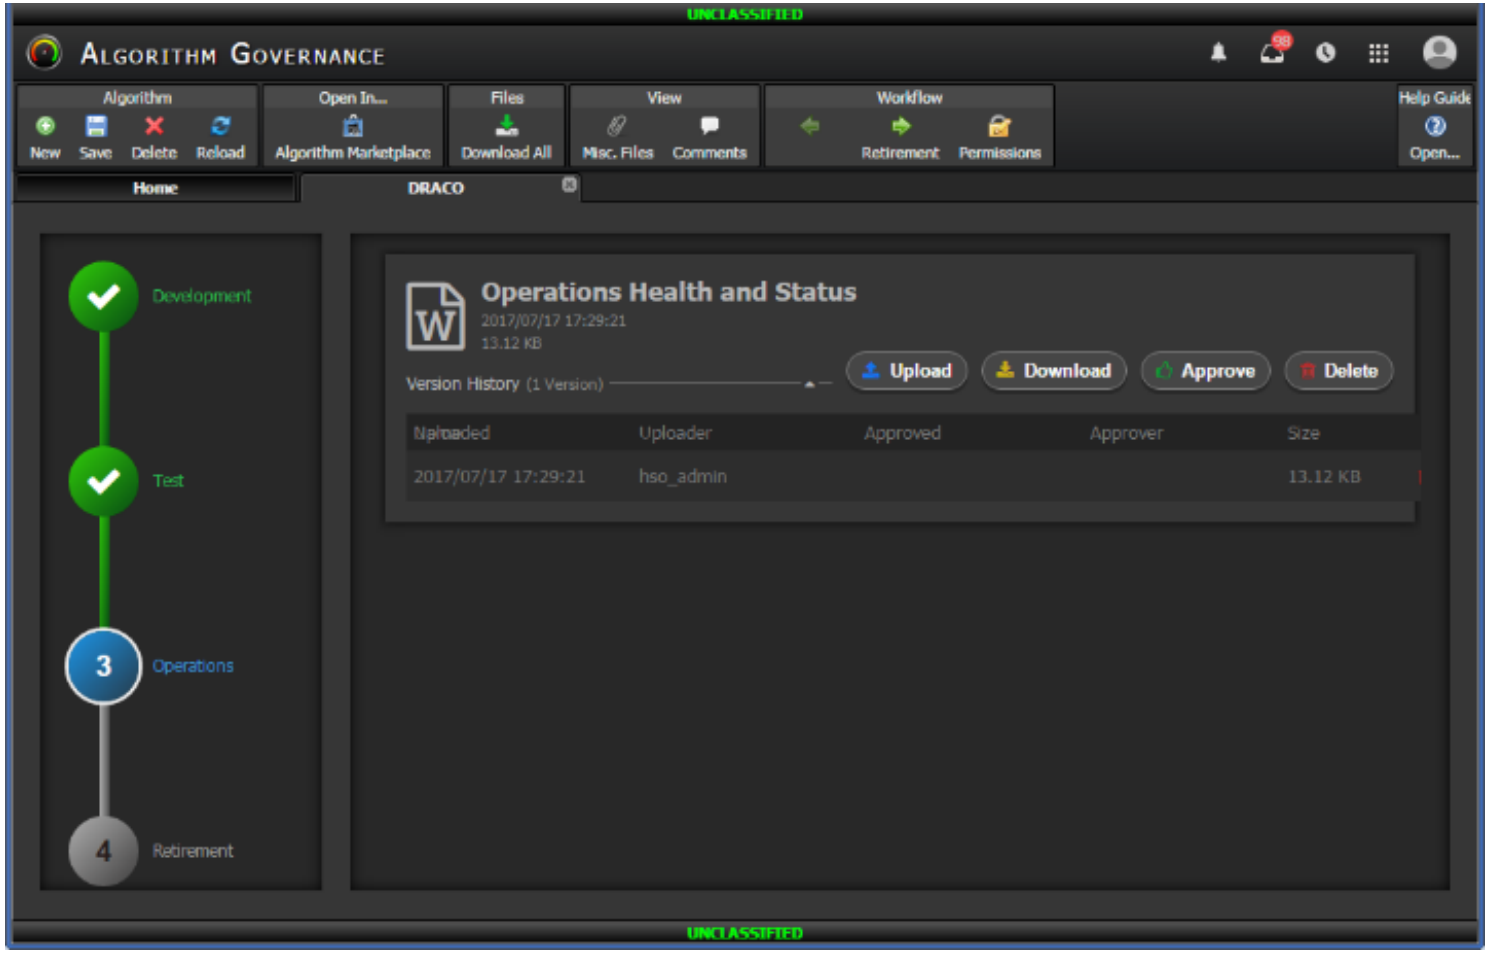

Algorithm Governance Workflow is based on Hydra's Workflow Service: flexible and extensible

Open APIs enable rapid integration with third-party clients and external systems

UNCLASSIFIED

REST API Tutorials Documentation

Product Search and Discovery

Based on geospatial information, what products are available?

What other searches are available for me to find?

If there aren"t existing searches, how can I create a search?

How do I find metadata for a product?

How do I find files associated with a product?

Deriving Analytics

What analytics can be derived from a given product?

How do I invoke analytical processing and place results in a desired location?

Product Tagging

How do I tag a product?

What products have a tag that I"m interested in?

How do I get all of the tags associated with a product?

Workflow and Product Management

What workflow queues are available?

What storage locations are available to place products?

UNCLASSIFIED

UNCLASSIFIED

REST API Tutorials Documentation

Algorithm Marketplace Service

The Algorithm Marketplace (AM) Service provides Algorithm Developers a location to publish their Algorithms to the community. By working with the HYDRA developer, they make their algorithms available to be invoked directly from the Algorithm Marketplace Application, or as part of an Automation Wizard in an automated fashion. The Algorithm Marketplace Service collaborates with: the Search Service to define a baseline search, which will identify good candidate products; the Algorithm Marketplace backend to store and run algorithms. Each algorithm in AM is associated with a base set of metadata constraints that limit the type of products that can be run through the algorithm; these constaints are defined in the baseline search. Examples of metadata constraints are as follows "modality is multispectral", "mission is Digital Globe", and "sun elevation angle is greater than 45 degrees". The intent of the constraints is to allow an algorithm to declare the subset of Hydra products that it can successfully operate on. Once an algorithm is published to the marketplace, it can be run on candidate products based on narrowband data or by viewing the thumbnail for the product, when available.

[BASE URL: /rest/am/v1 , API VERSION: 1.0]

Algorithms

Show/Hide | List Operations | Expand Operations | Raw

GET	/algorithms	Get Algorithms
GET	/algorithms/simple	Get Algorithms without Parameters
GET	/algorithms/{algorithmName}	Get Algorithm by Algorithm Name
GET	/algorithms/modality/{modality}	Get Algorithms by Modality

Parameters

Show/Hide | List Operations | Expand Operations | Raw

GET	/parameters/{algorithmName}	Get Parameters by Algorithm Name

UNCLASSIFIED

Self-describing APIs reduce the burden to generate large ICD documentation

DAGR (Distributed All-source GEOINT-analytics Resource)

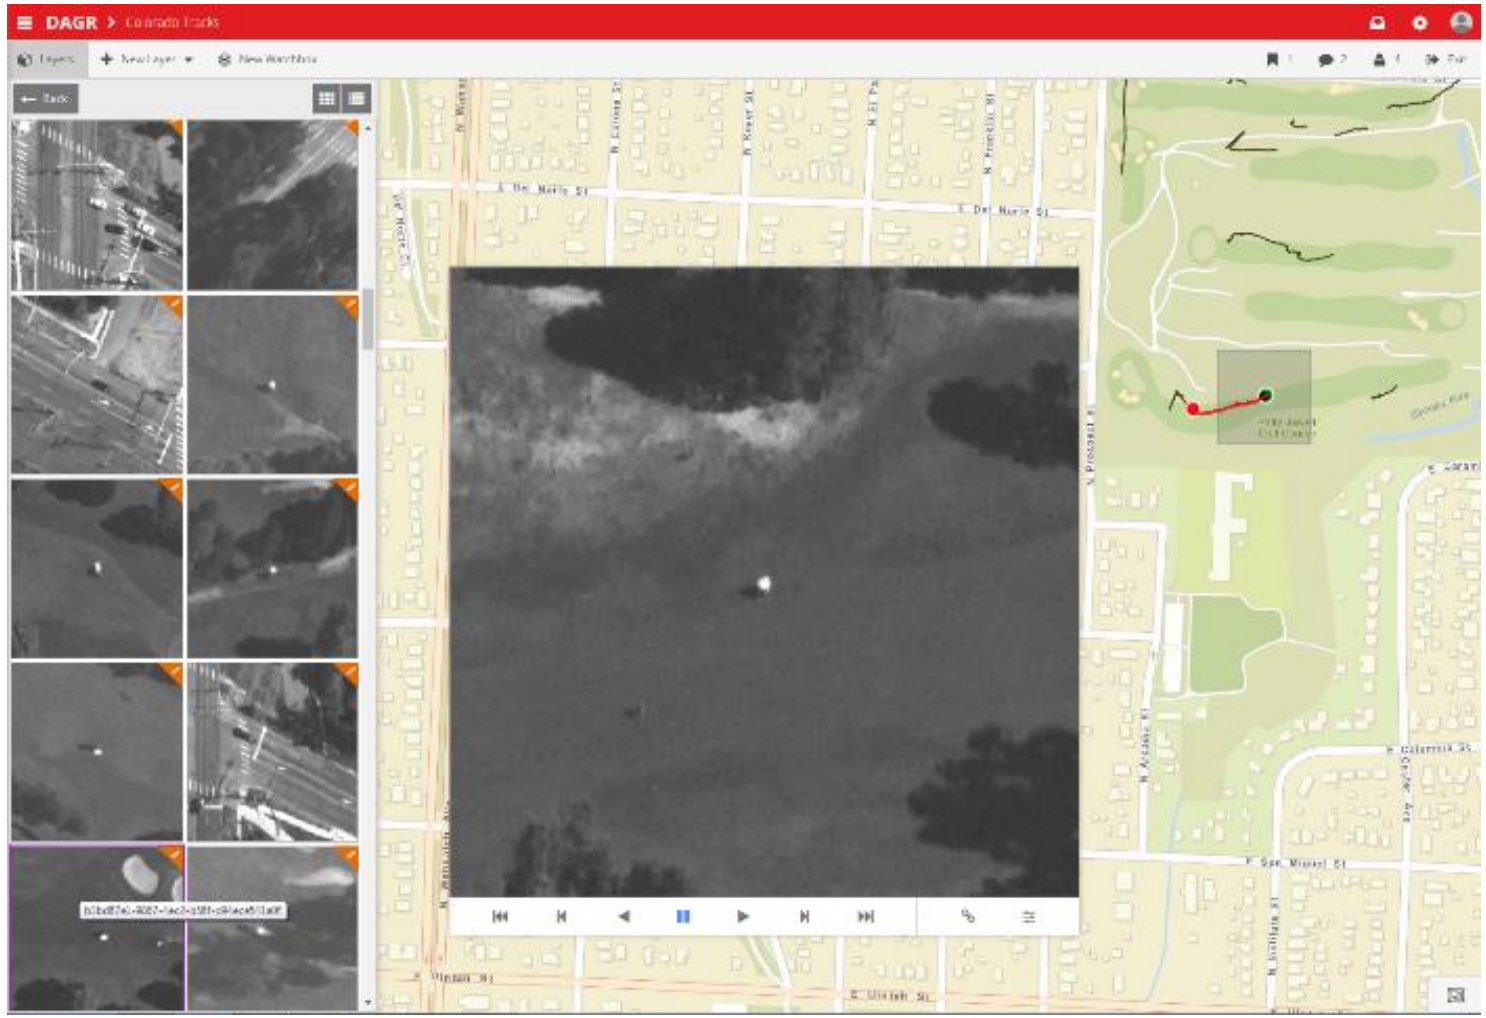

DAGR modernizes the analyst workforce with automated multi-intelligence workflows and workspace collaboration to solve complex intelligence questions

Map-based analytics recommendation

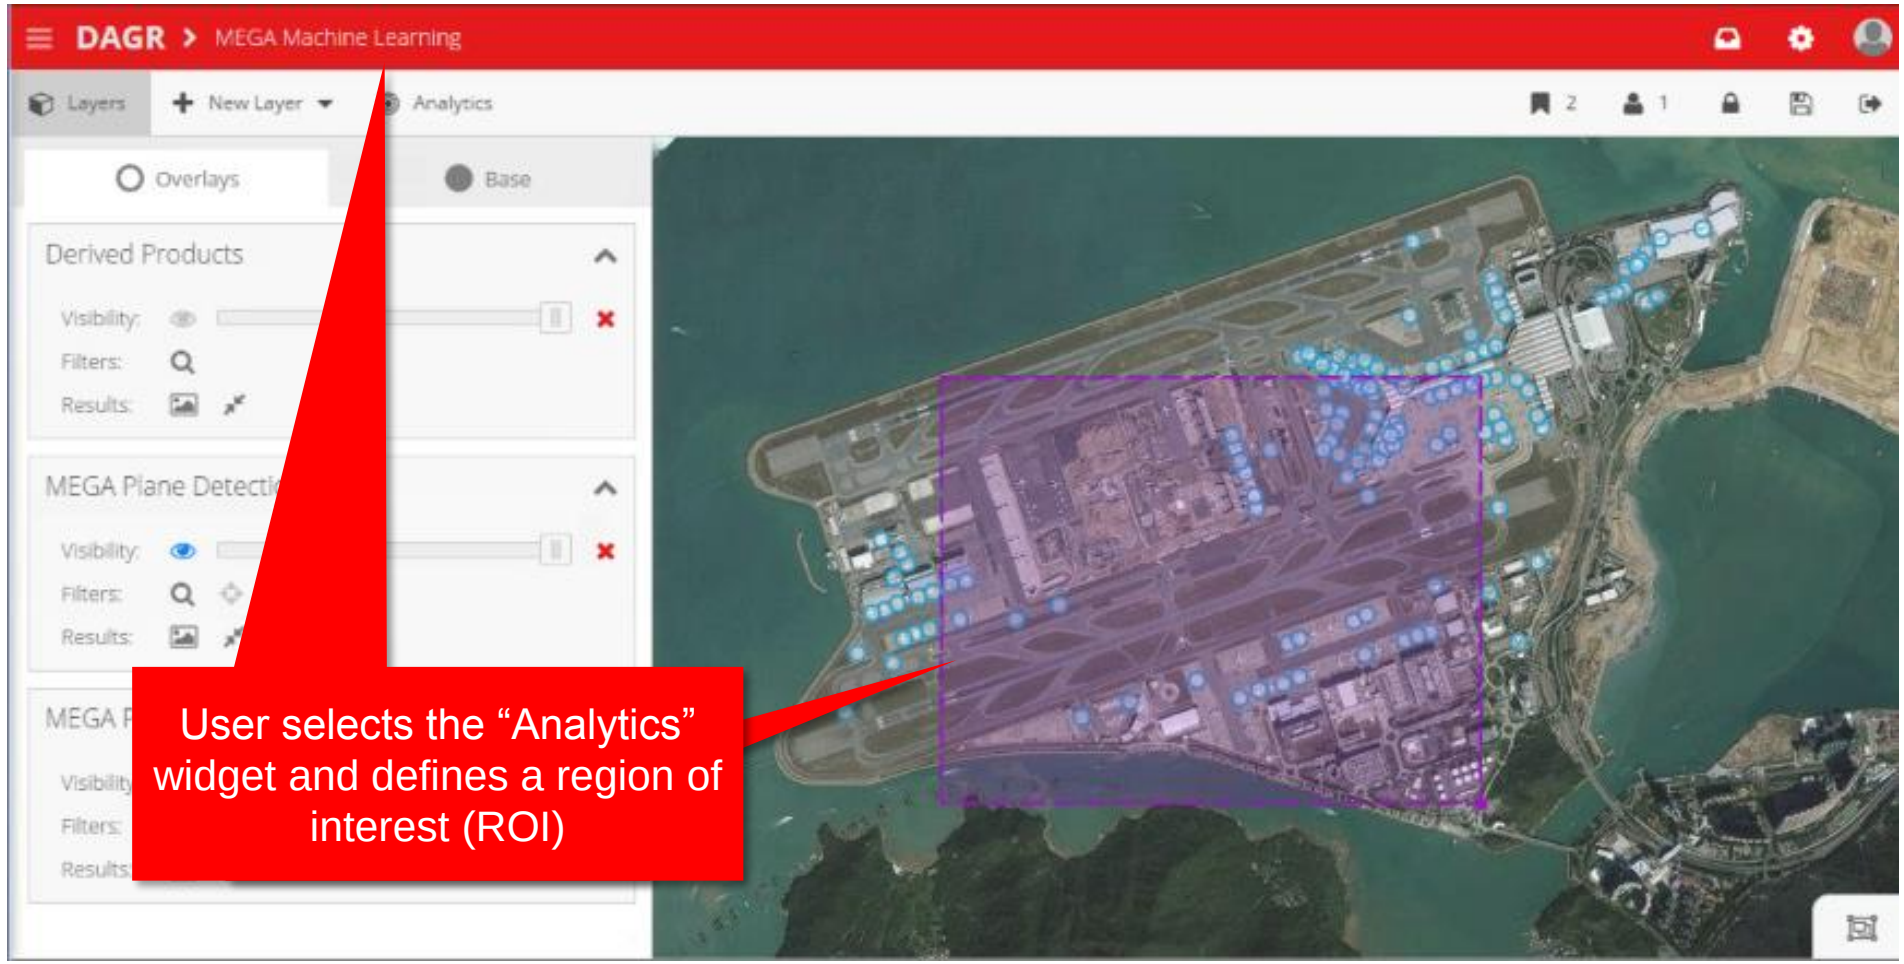

Map-based analytics recommendation available via Hydra-API for third-party client integration

Map-based analytics recommendation

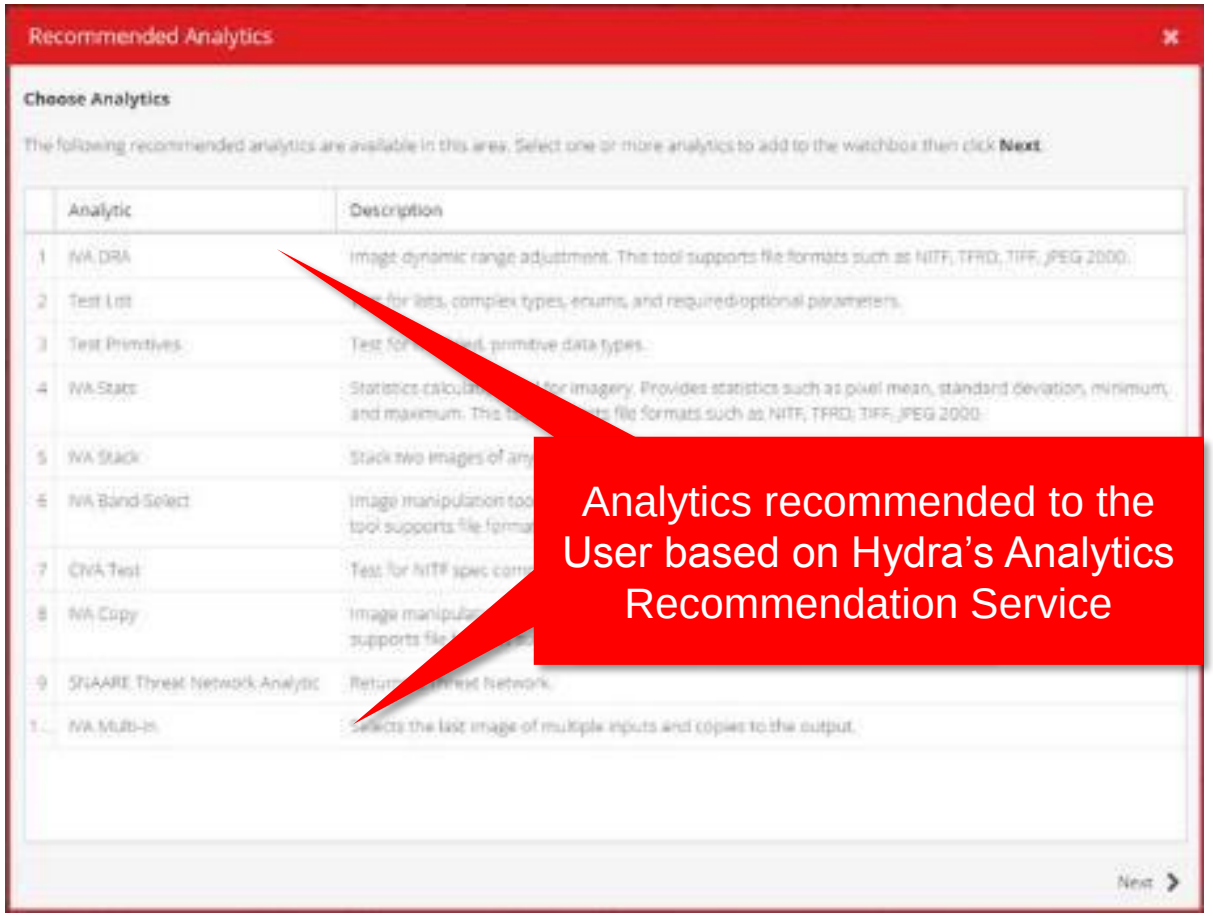

Analytics recommended to the User based on Hydra's Analytics Recommendation Service

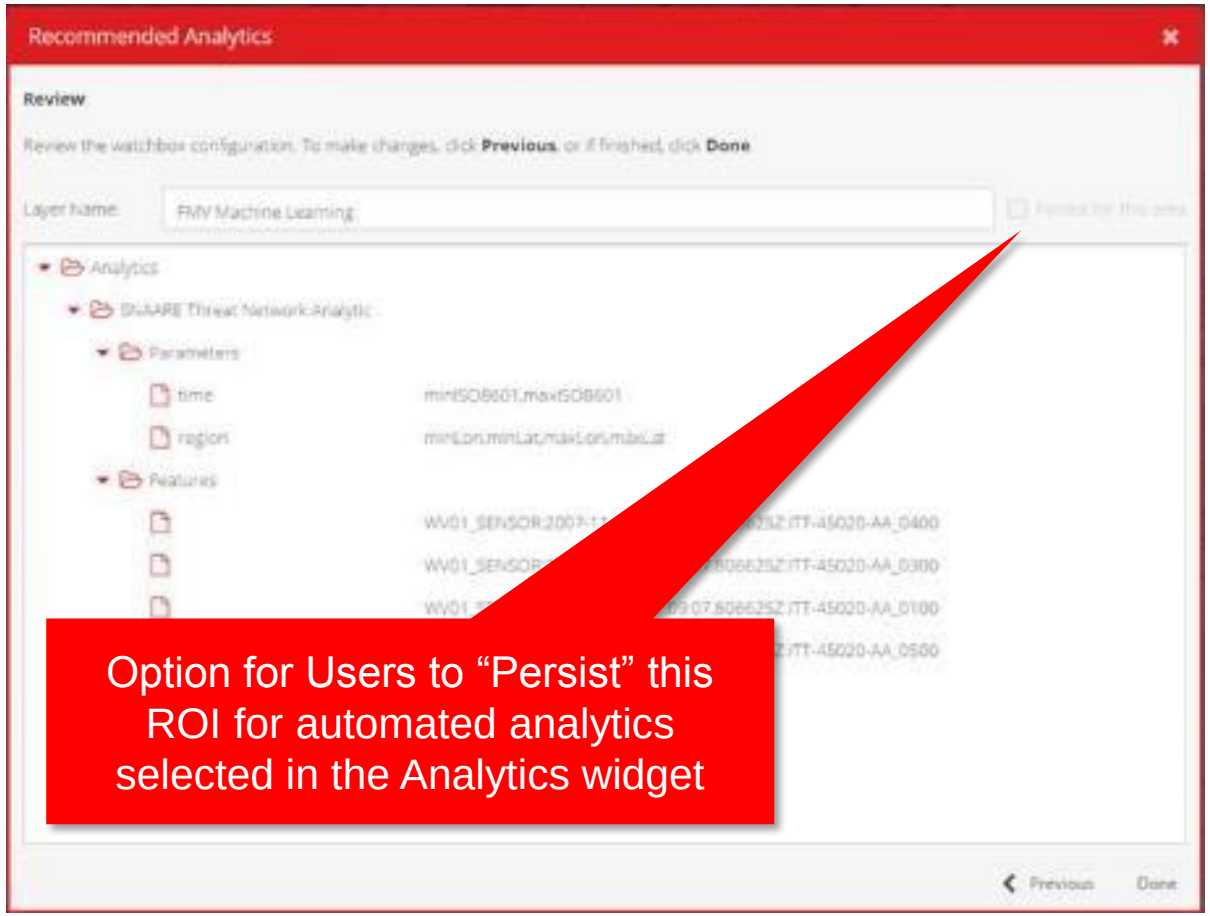

Option for Users to "Persist" this ROI for automated analytics selected in the Analytics widget

Hydra-API available to automate analytics based on ROI

Hydra's flexible cataloging supports machine learning algorithms

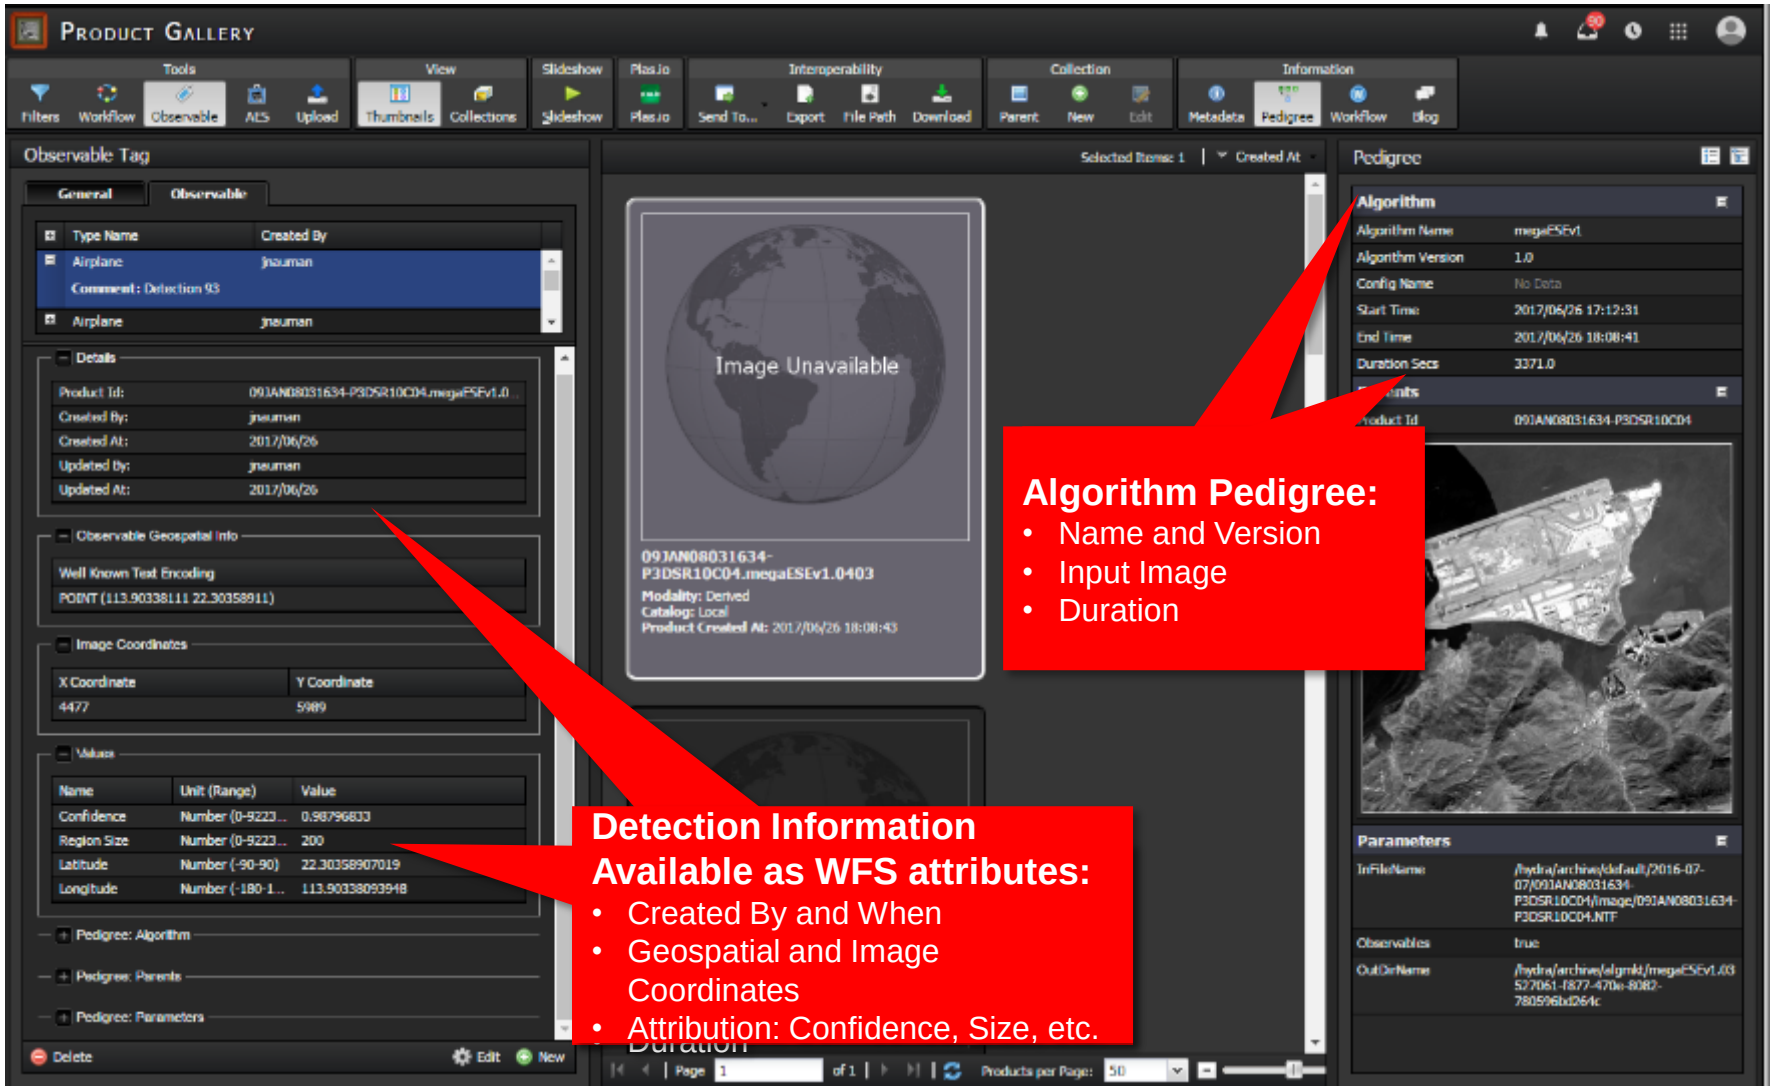

Hydra Flexible Cataloguing "Standard" Enables Rapid Insertion of Developer-defined Algorithm Attributes/Observables

Google Earth visualization of machine learning output

Standards-based discovery, dissemination and rapid access

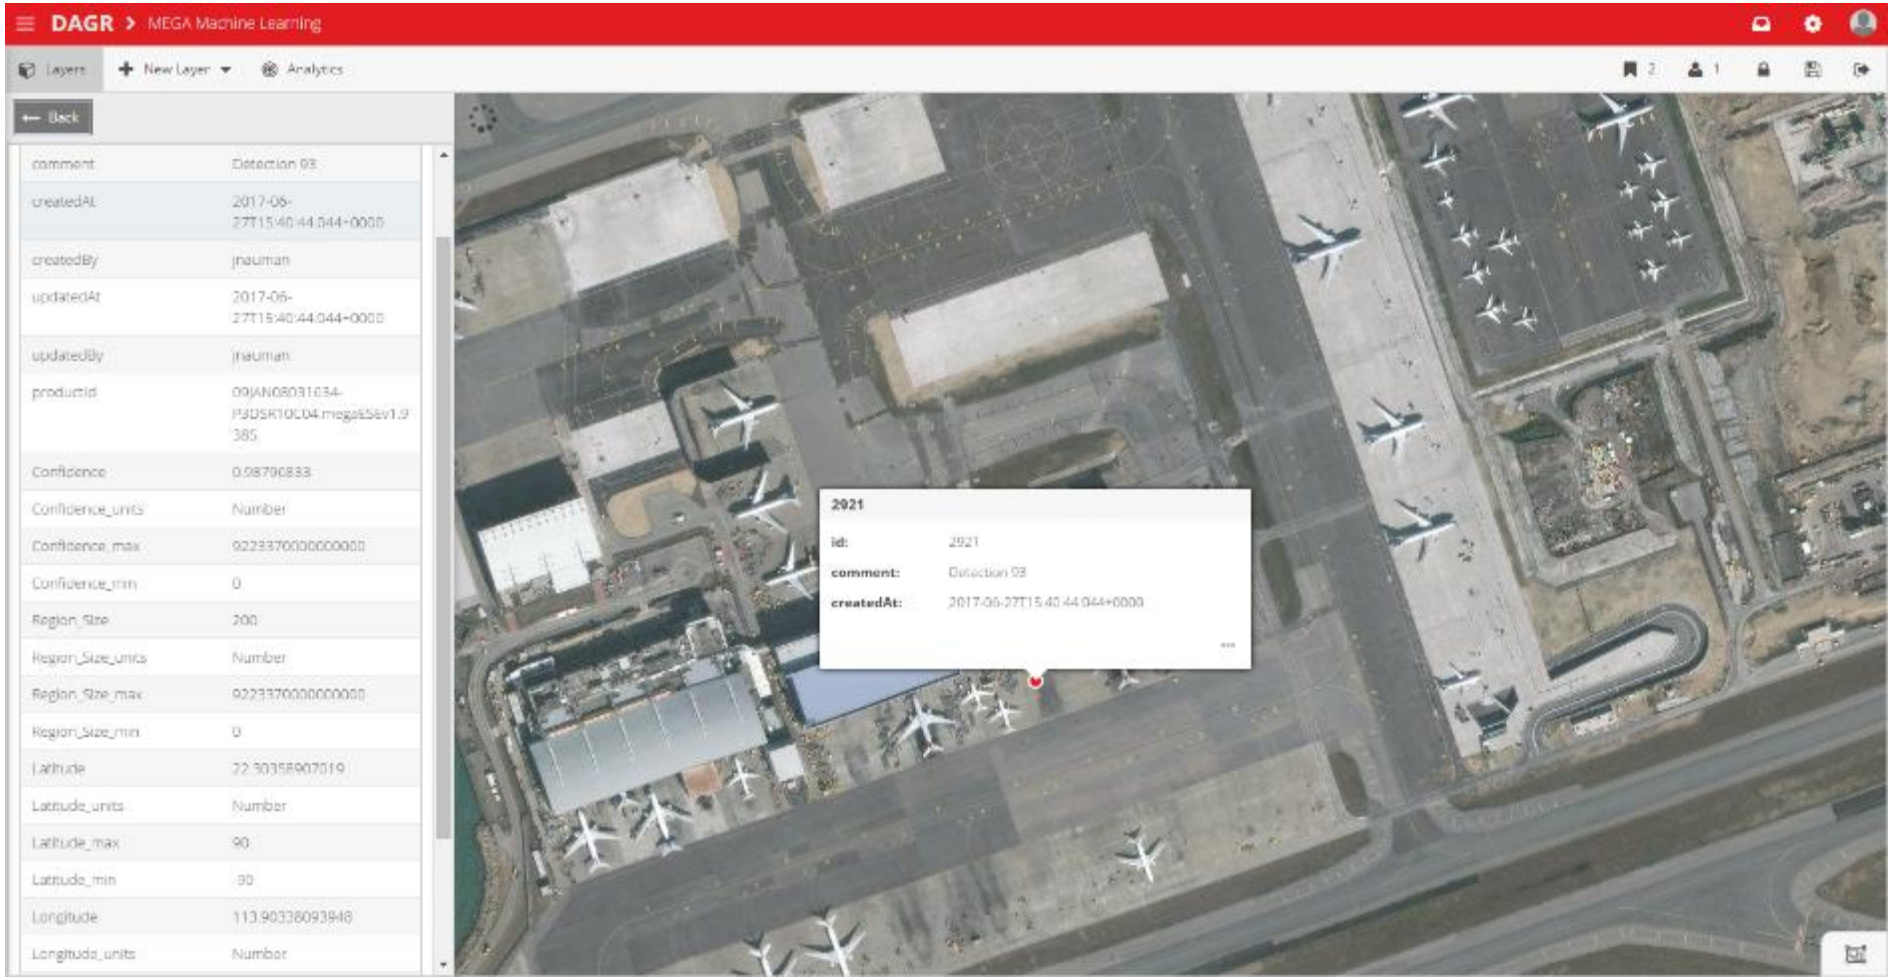

Hydra exposes detections as WFS, rapidly consumable by third-party clients

Geometric change detection of multi-source point clouds

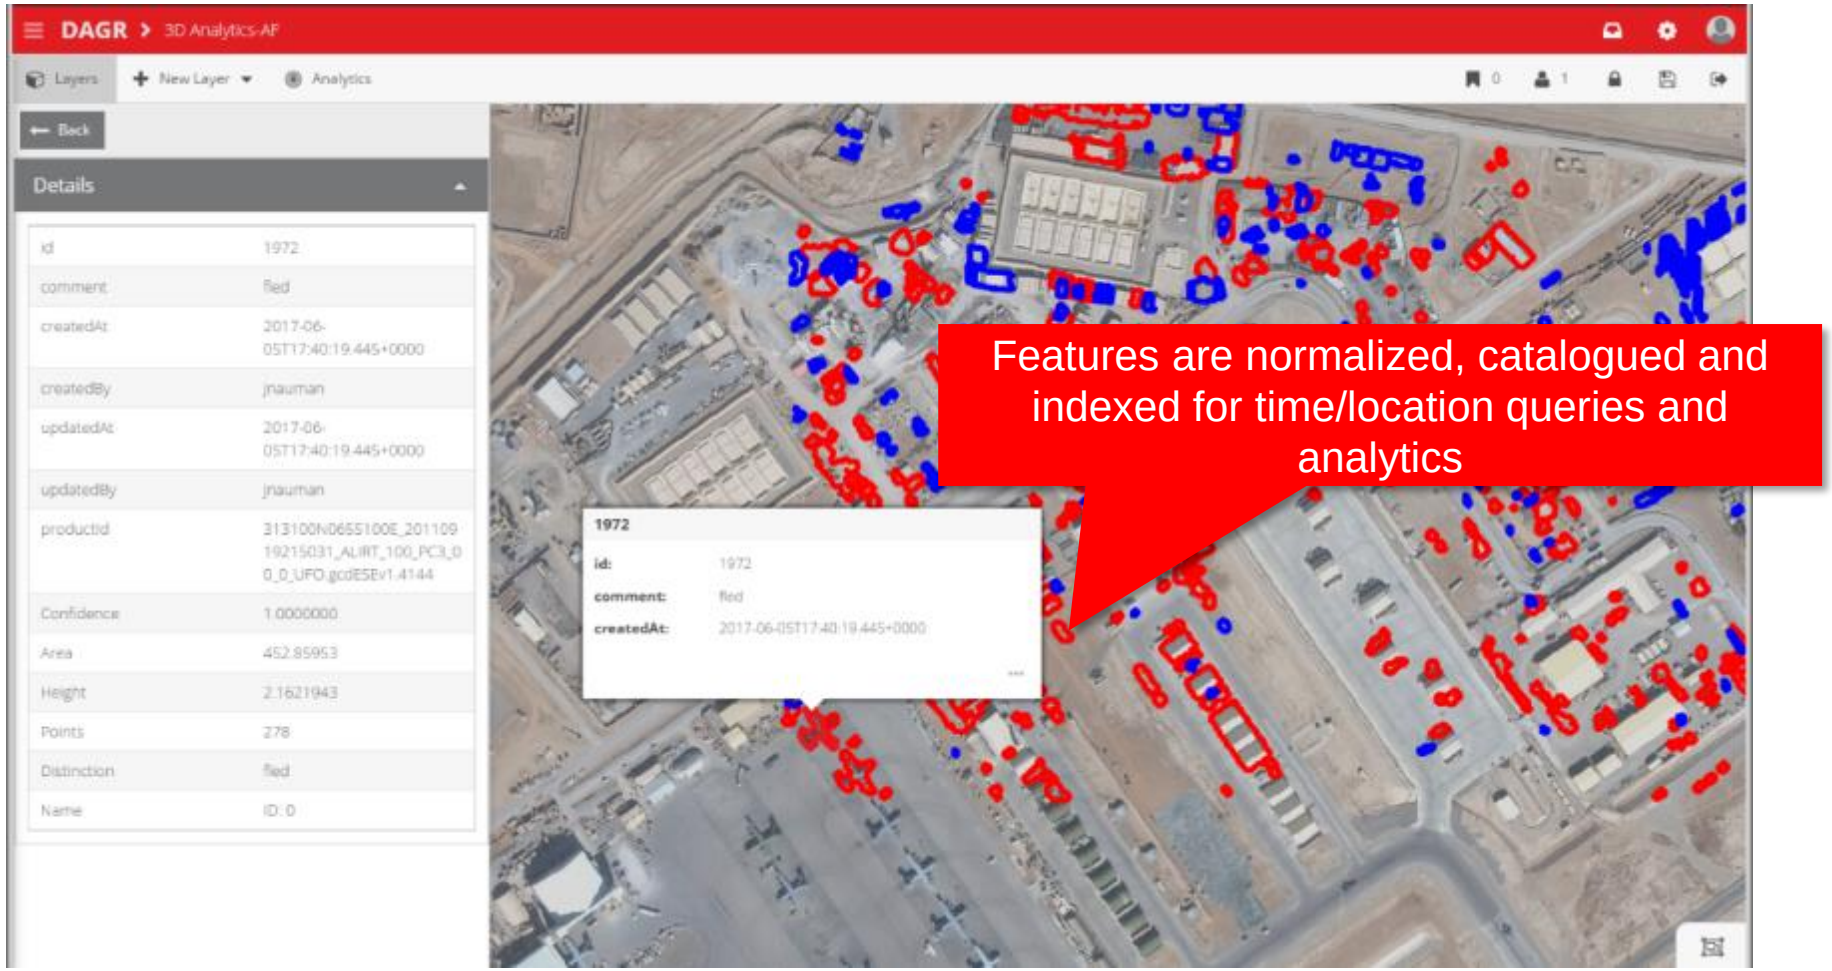

Hydra enables an Integrated Data Layer (IDL) composed of input sources, detected objects/signatures/detections, and multi-INT correlation analytics

Hydra applications
Closed-loop MEGA training via Hydra