Defence AI Centre
Ministry of Defence

The Uncertain Object

Application of Conformal Prediction to Aerial and Satellite Images

Vicky Copley, Greg Finlay, Ben Hiett

DSTL DAIC-X Applied V^3 Project

DSTL/TR161984

Strategic Command
Defence Digital

dstl

de&s

Outline

- Setting
- Data
- Methods
- Detections
- Localisation
- Summary and next steps

Setting

Intelligence, Surveillance and Reconnaissance (ISR)

- Intelligence generation, surveillance and reconnaissance are a centuries old part of military operations
- Modern sensors mean that data is no longer scarce
- Challenge to convert data to useable intelligence
  - making best use of limited compute resources and human bandwidth
  - providing performance guarantees for underlying machine learning models, enabling trust


1


2

1 Bundesarchiv_Bild_183-R01996,_Brieftaube_mit_Fotokamera.jpg: o.Ang.derivative work: Hans Adler (talk) - Bundesarchiv_Bild_183-R01996,_Brieftaube_mit_Fotokamera.jpg, CC BY-SA 3.0 de, https://commons.wikimedia.org/w/index.php?curid=7163150
2 Shoka at the English-language Wikipedia, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=16296864

Intelligence

- Making sense of raw data, potentially from multiple sources
    - intelligence analysts are typically looking for ‘needles in haystacks’
    - possible representations of a 100 piece jigsaw from 7 pieces
- Probability-based objective judgements are core to intelligence assessments
    - intelligence probability yardstick
    - consistent terminology

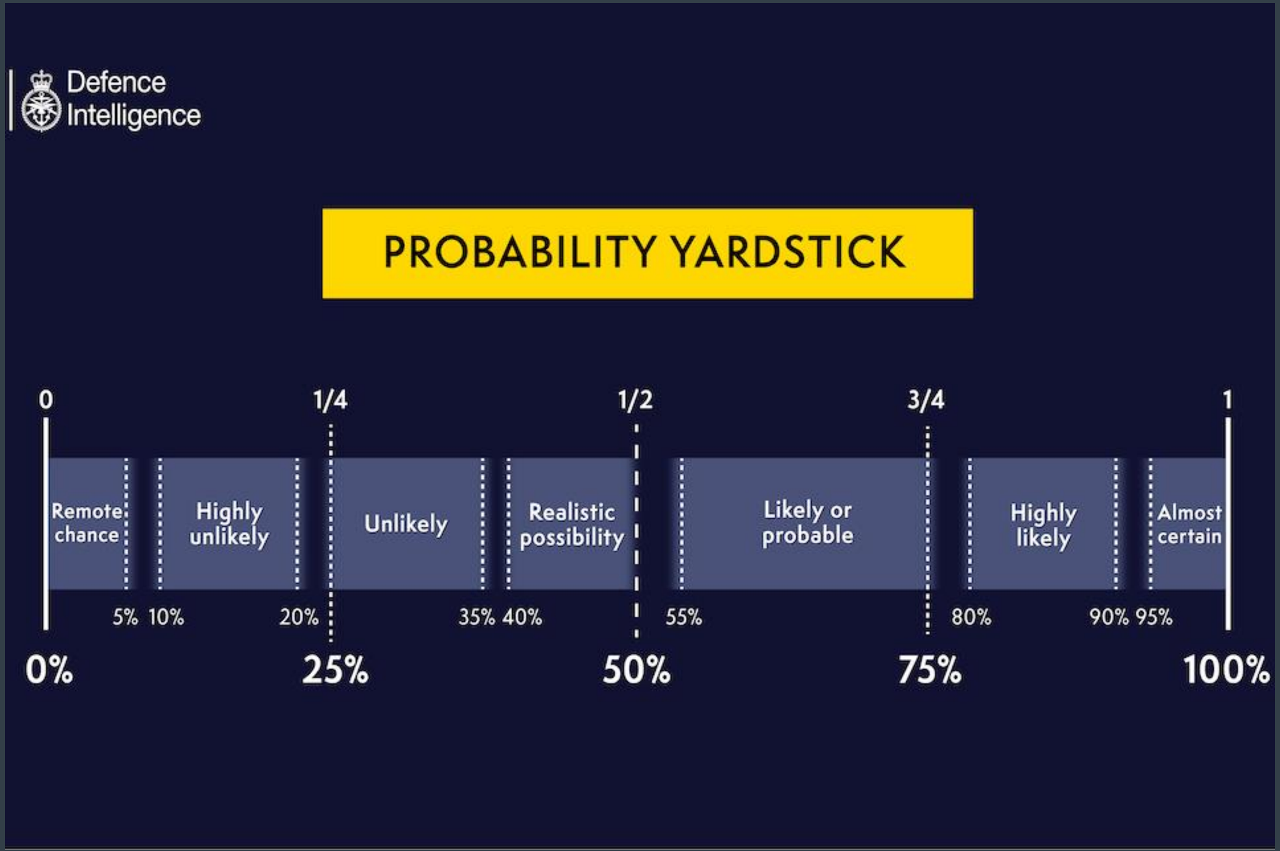

Source: https://www.gov.uk/government/news/defence-intelligence-communicating-probability

Motivation - Tyche satellite

- Small 150kg satellite launched 16 August 2024 by UK Space Command
  - space-based ISR
  - able to capture daytime images and videos of the Earth's surface
  - in-orbit processing, limited bandwidth (edge)
- How can we best quantify the uncertainty in object detection models to provide key information to intelligence analysts?
  - performance guarantees for model predictions
  - optimise use of limited compute and bandwidth

Aims

- Use conformal prediction approaches to quantify (some of) the uncertainty which exists in trained object detection models
- Examine various loss functions and nonconformity measures
- Assess performance using satellite and aerial remotely sensed data typical of an ISR setting

Data

Remotely sensed data

- 3.5m resolution multispectral satellite data[1]

- Drone data of Aerial Floating Objects (AFO)[2]

1 Vision-1 data courtesy of Airbus Intelligence UK and UK MOD ARTEMIS programme
2 https://www.kaggle.com/datasets/jangsienicajzkowy/afo-aerial-dataset-of-floating-objects, used under CC licence https://creativecommons.org/licenses/by-nc-sa/3.0/igo/

Methods

Object detection

- Subfield of Computer Vision (CV)
  - locates instances of objects in images or videos and draws bounding boxes around them
  - fitted using convolutional neural networks
  - many off-the-shelf model architectures
- Yolo v5 nano architecture
  - requires less processing resource
  - suitable for use on lightweight ISR satellites
  - two trained models, one for each data source

Source: https://uk.mathworks.com/discovery/object-detection.html

Uncertainty in object detection

- At least 6 sources of uncertainty in an object detection model[1]
  - missed detections
  - spurious detections
  - localisation
  - …
- Non-conformity measures in Yolo model
  - missed detections/spurious detections:
    - object score
    - confidence score (object score*maximum class probability)
  - localisation
    - intersection over union

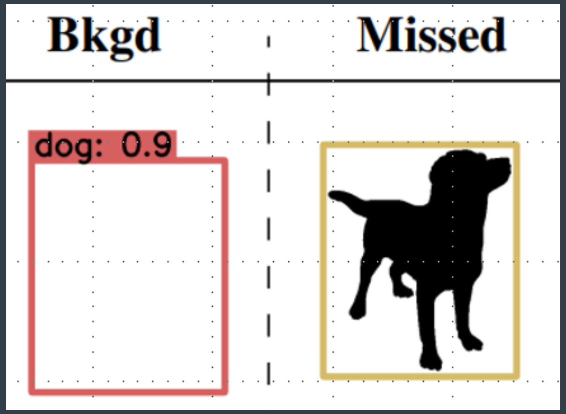

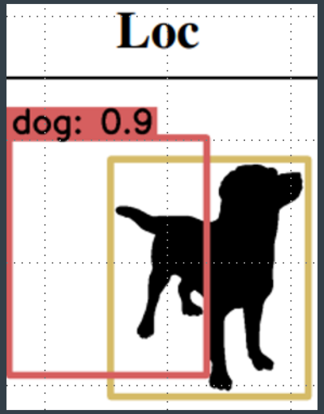

[1] and images source: TIDE: A General Toolbox for Identifying Object Detection Errors
https://arxiv.org/pdf/2008.08115.pdf

Raw Yolo output

- In inference mode Yolo initially predicts thousands of boxes per image
  - uses input parameters to fine tune the raw output, with default settings
  - consider that initial boxes include all objects of interest
  - use conformal risk control to parameterise inference run time input parameters, providing theoretical guarantees

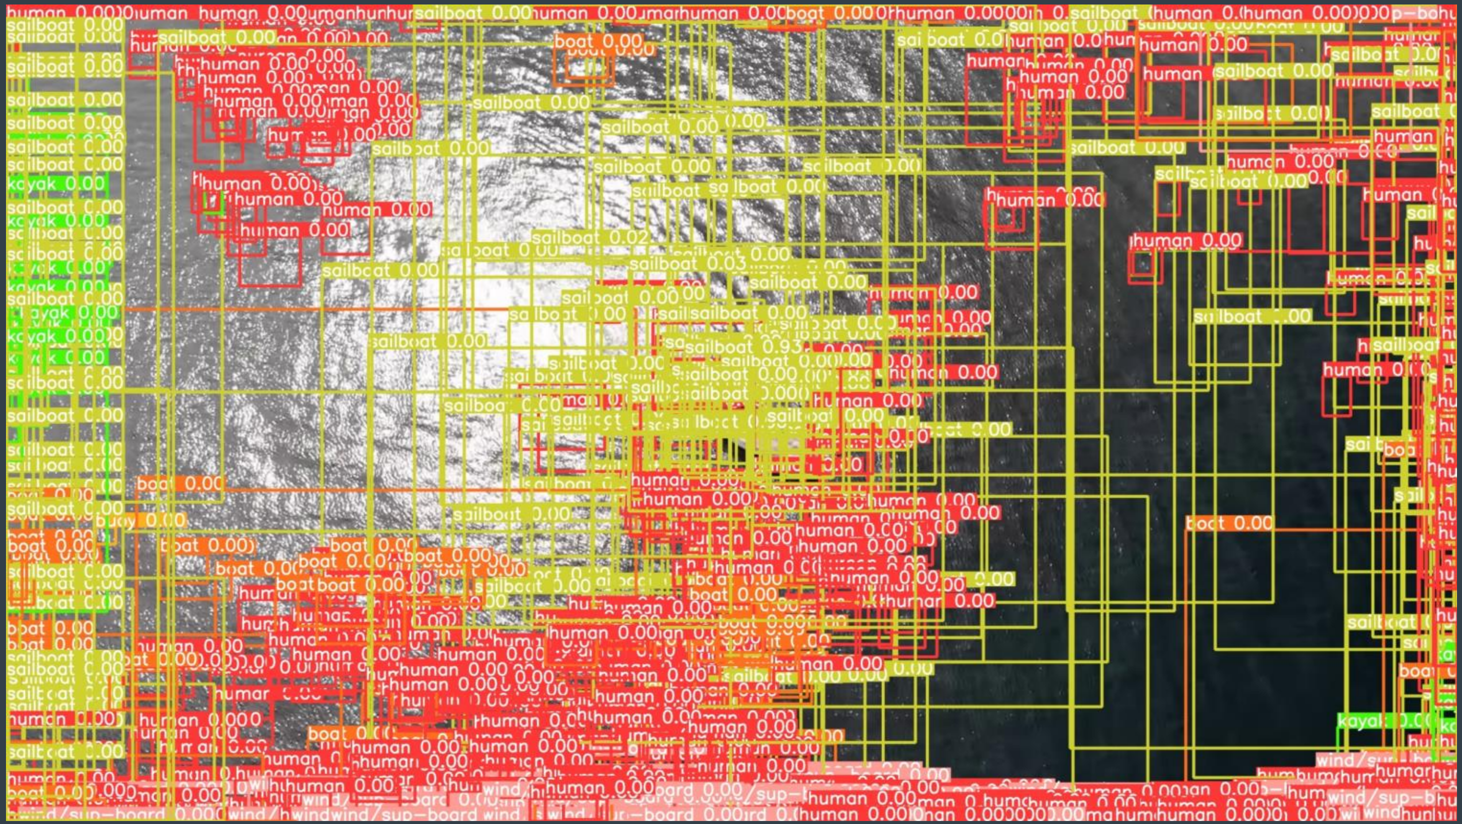

Ground truth

Conformal risk control

- Method to control other risk functions as well as coverage[1]
  - generalises conformal prediction to provide guarantees with the form

$$\mathbb{E}\left[\ell\left(\mathcal{C}(X_{n+1}), Y_{n+1}\right)\right] \leq \alpha$$

  - for any bounded loss function ℓ that shrinks as $\mathcal{C}(X_{n+1})$ grows
    - $\mathcal{C}$ is a function of the model and calibration data that outputs a prediction set
    - α is a user-specified error rate
- Miscoverage and false negative rate (FNR) loss functions considered for detection problem

[1] A. N. Angelopoulos, S. Bates, A. Fisch, L. Lei, T. Schuster. Conformal Risk Control. 2022. https://arxiv.org/abs/2208.02814

Experimental design for detections

- Various simulation experiments:

Non-conformity measure	Model	Alpha	Performance metrics
1 - confidence score 1 - object score	Vision-1	0.5	Miscoverage False negative rate Detection multiple
		0.4	
		0.3	
1 - confidence score 1 - object score	AFO	0.3	Miscoverage False negative rate Detection multiple
		0.2	
		0.1	

- false negatives are generally considered to be the costliest type of error in an intelligence context …
  - also need to avoid so many false positives that the analyst becomes frustrated with the system[1]

[1] Knack et al. 2022. Human Machine Teaming in Intelligence Analysis. Alan Turing Institute Research Report https://cetas.turing.ac.uk/publications/human-machine-teaming-intelligence-analysis

Results (Detections)

Distribution of scores by model data

- Charts show the distribution and cumulative distribution of confidence scores matched to ground truth boxes

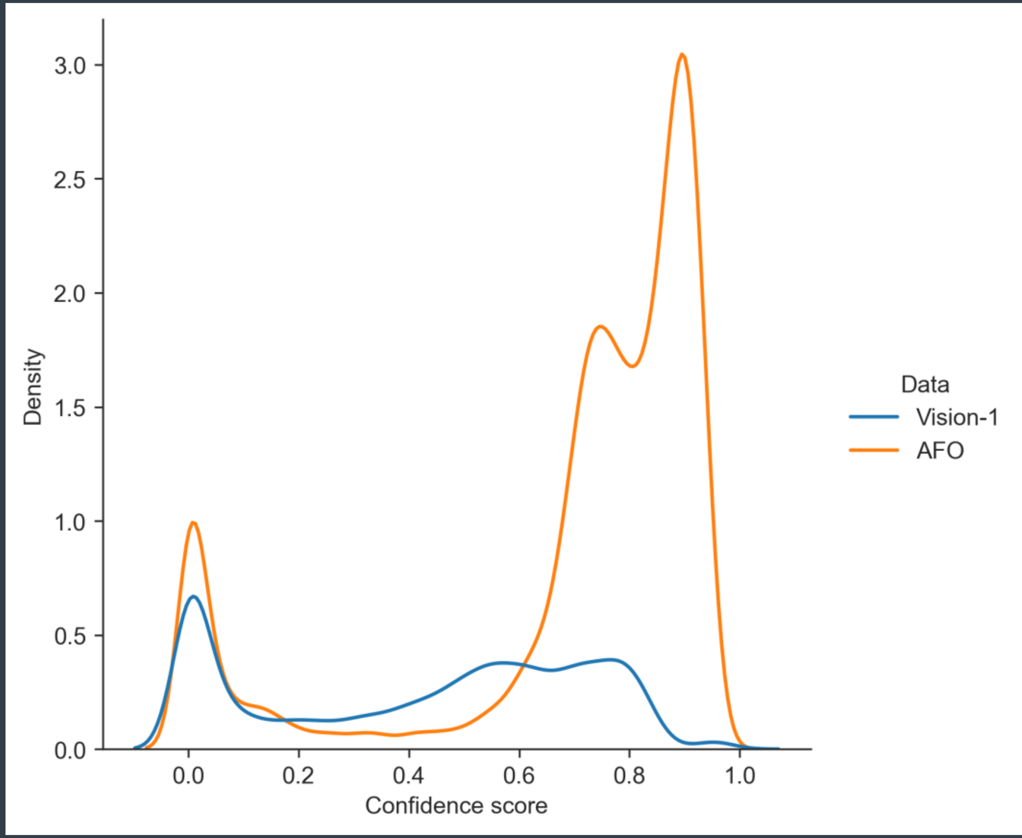

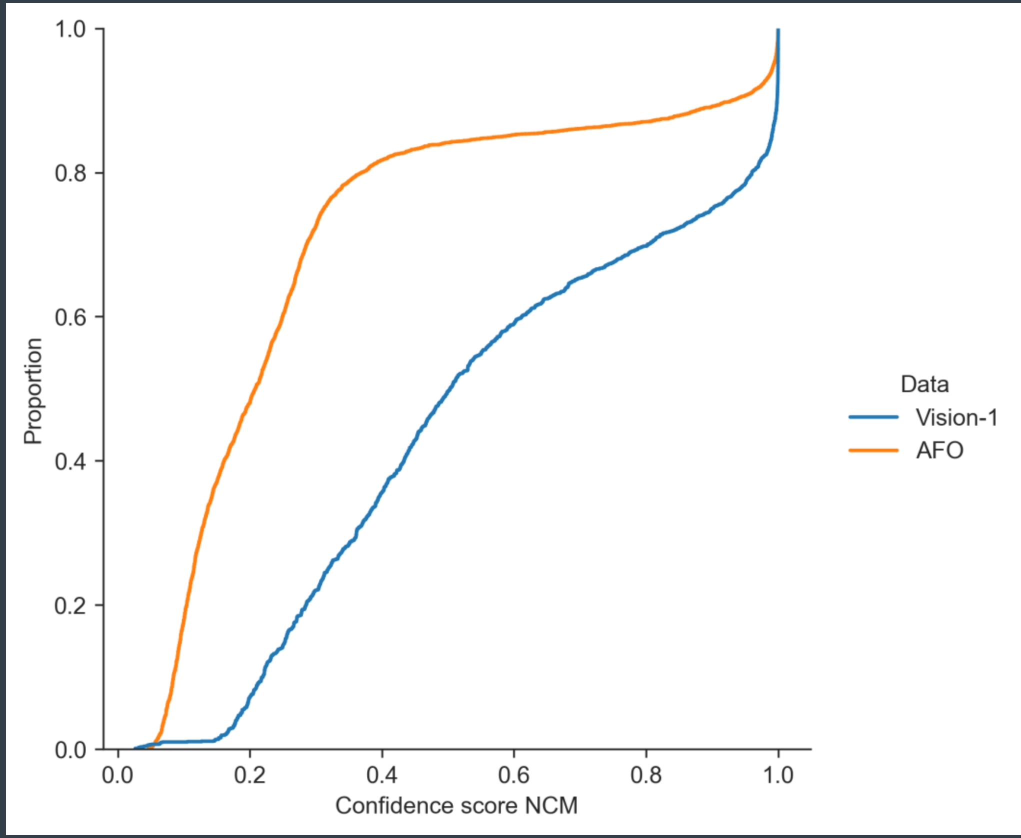

Vision-1 confidence score FNR results

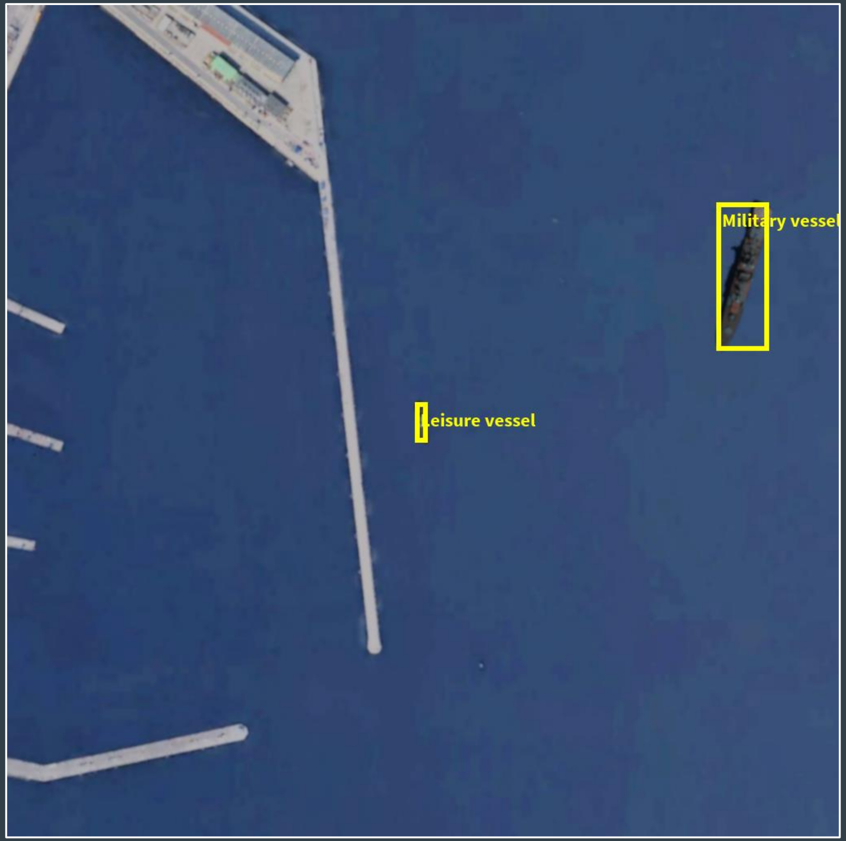

Ground truth

α = 0.5

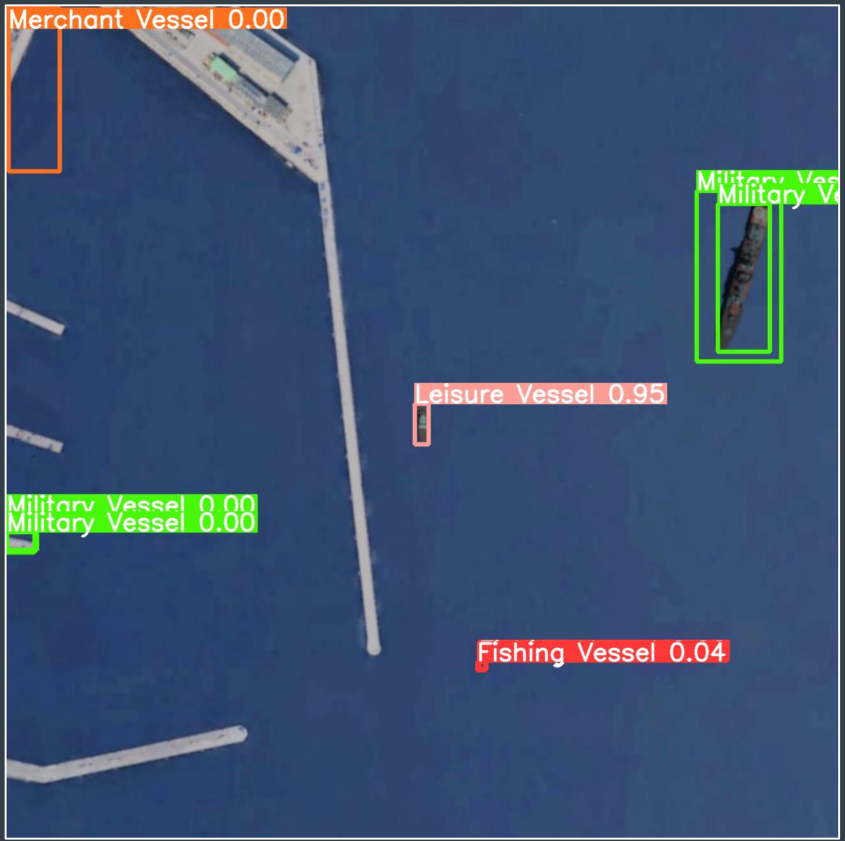

α = 0.3

Localisation

Bounding box approach

- Applies conformal to the bounding boxes themselves, not the prediction
- Gives coverage guarantees on whether the predicted bounding box contains the true object within it
- Useful to an intelligence analyst who needs to know they are not missing parts of a detected object

Performance metrics

- For localisation application we used Intersection over Union and Bounding Box Area ratio
- Average % of the original IoU that is retained and average increase in bounding box area, post conformal prediction
- Smaller values of α generally resulted in larger bounding boxes, thus less precise predictions

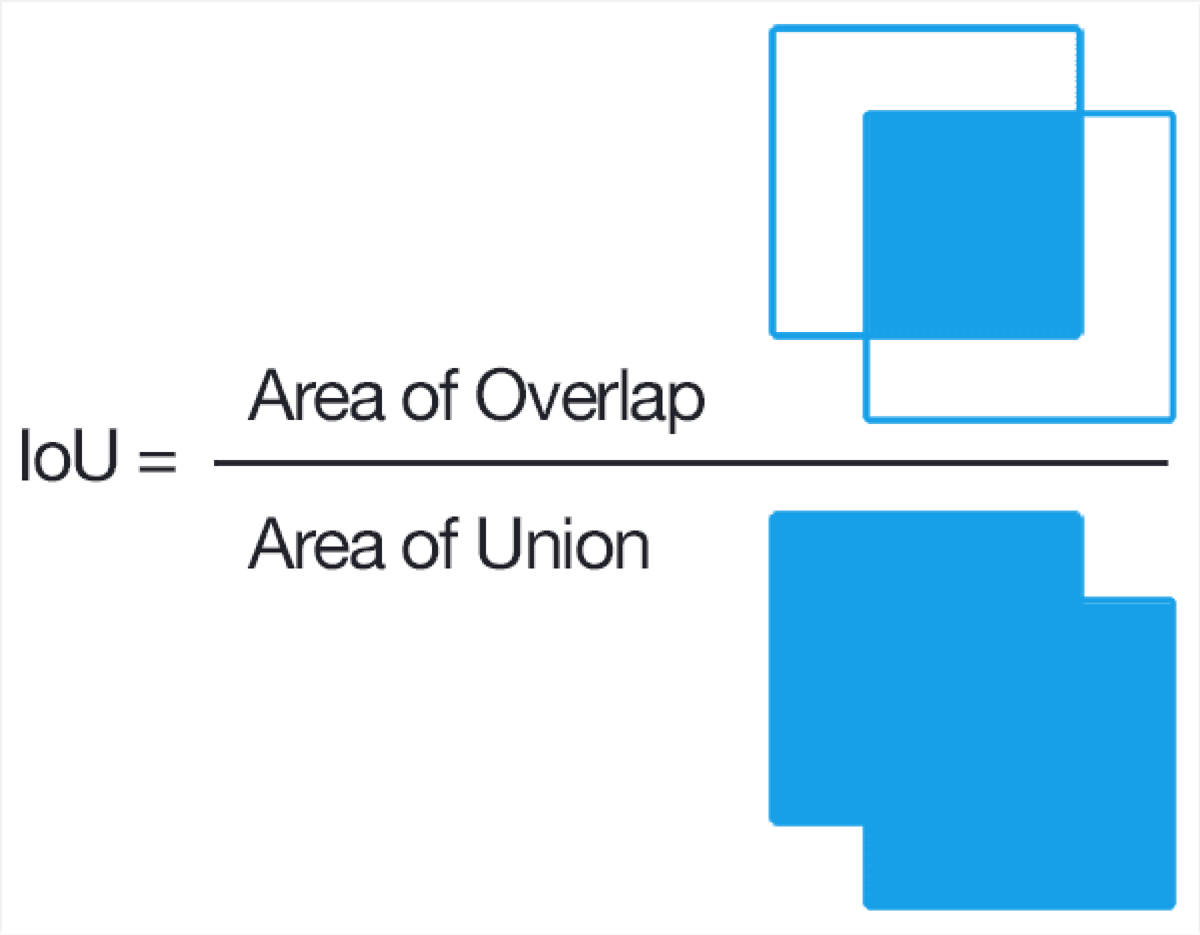

Image from https://pyimagesearch.com/2016/11/07/intersection-over-union-iou-for-object-detection/

Bounding box example

- Applied to Vision-1 model
  - providing object level coverage guarantee on the chance a predicted bounding box will contain the true object within it
- α = 0.2
- In many cases did not appreciably change the size of bounding boxes
- But a poor predictor can have a significant affect on size of boxes

Localisation results

NCM	α	Coverage	Average IOU retained	Average box area change
Bounding box difference quantile	0.4	0.8222	0.6593	1.848
	0.3	0.8444	0.6399	1.908
	0.2	0.9171	0.5688	2.217
	0.1	0.9436	0.5254	2.423
Conservative bounding box difference quantile	0.4	0.9291	0.5865	2.152
	0.3	0.9291	0.5705	2.225
	0.2	0.9675	0.4895	2.665
	0.1	0.9709	0.4567	2.849

Summary

Summary of findings

- Object detection with remotely sensed data can be challenging
  - predicted scores of some ground truth objects clustered around zero in both models
  - distribution of non-conformity scores very different between models
- Conformal approaches can be used to pick Yolo command line parameter settings which guarantee (low) error rates in detections
  - exposing model uncertainty in a manner which is lightweight, flexible and versatile
  - forming part of verification and validation procedures for approving and releasing a model into operational use
  - filtering out extraneous information and providing analyst with a level of confidence in the output

Future work

- Conformal prediction for object detection ‘as-a-tool’
  - Use within an overarching assurance test framework
  - Plug & Play ‘tool’ within a larger ‘toolkit’
  - Working toward ‘operational deployment’ of the technique
  - In pursuit of ‘safe and dependable AI’ for MoD
- Conformal prediction for Large Language Models
- Conformal prediction in ML Training – calibration of model consistency classifier to accelerate training
- Monitoring data drift in satellite data to prompt retraining

Questions