FRONTIER SI

AUSTRALIAN GEOSPATIAL-INTELLIGENCE ORGANISATION (AGO) ANALYTICS LAB PROGRAM

ROUND 2

PROPOSAL BRIEFING DOCUMENT

Laura Spelbrink, Program Manager
lspelbrink@frontiersi.com.au

Phil Delaney, Chief Innovation and Delivery Officer
pdelaney@frontiersi.com.au

Document Control
Date: July 2021
Version: 1.6
Reviewed By: Laura Spelbrink
Approved By: Phil Delaney

Objective

The AGO's Analytics Labs Program (AGO Labs), coordinated through FrontierSI, was created to help AGO better access innovative technology and trial new ways of working with industry. Specifically, AGO is keen to attract a wider pool of companies and technologies to draw on for automated geospatial intelligence. The Program aims to address AGO challenges and build new industry capability through short-term projects with Australian and New Zealand-based companies. AGO Labs Challenge Topics focus on machine learning and analytics challenges for producing automated imagery analysis.

Successful demonstrator projects may have the opportunity to progress to an additional limited operational testing phase.

What is AGO Labs?

AGO Labs is a way for AGO to rapidly assess new technology capability through a challenge-based innovation program. From this process, learnings can be gathered about the barriers and opportunities for AGO to work with an increased breadth of companies. In turn, this could lead to a long-term mechanism for AGO to access, and provide, a pipeline of activities to industry for testing new innovations and thinking in analytics.

Initial successful projects will be funded up to $100,000 for Proof-of-Concept projects for projects up to six months in length. An additional $50,000 and 3 months may be provided to some projects to undertake limited operational testing, though this process will be considered as a stage gate at the end of the Proof of Concept, and the focus on the Call for Proposals should focus on the Proof-of-Concept stage, rather than the operational testing stage.

There will be two rounds of AGO Labs proposal calls. One in November 2020 and one in the June 2021. Challenge topics listed herein are for Round 2.

The four AGO Capability Challenge Topics for Round 2 are briefly described below. It is recommended that organisations read the full description of each challenge topic to understand background, use cases and evaluation criteria for each of the challenge topics. Further information is given in Appendix A of this document.

Round 2 Challenge Topics

Challenge Topic 4: What Vessel is That?

This challenge aims to investigate whether algorithms can be run across commercial Synthetic Aperture Radar (SAR) satellite data to classify maritime vessels.

Challenge Topic 5: Made You Look

This topic is a research-focussed investigation on the possibility of encountering and detecting malicious modification of commercial imagery products.

Challenge Topic 6: What's Behind that Hill?

This challenge aims to determine the impact of masking of imagery by terrain or structures on automated object detection.

Challenge Topic 7: Detections from Geometry

This challenge aims to investigate whether machine learning techniques can identify categories of objects/facilities of interest based on the input of known dimensions and/or shapes.

Timeline

Interested parties may submit an AGO Labs project proposal by completing a short proposal, using the template downloaded from the AGO Labs page, and emailing to the FrontierSI AGO Labs Project Manager, Laura Spelbrink, at agolabs@frontiersi.com.au **by 5:00 pm (AEDST), 3 August 2021**.

Projects will be shortlisted by a panel with representatives from both AGO and FrontierSI. Applicants may be contacted to provided further details on their proposals.

Successful Projects will be notified in August 2021. Projects are expected to start in September 2021 and be delivered within a timeframe of six months.

Budget

Proof-of-Concept projects will be funded with a budget of up to $100,000. There is no requirement that additional funding will be provided to the project by applicants.

As this activity will help companies pilot technology with AGO for potential future deployment, it is expected that applicants will not operate with full commercial rates, but instead will budget the project at-cost plus 30% overheads.

For projects that are considered particularly successful and make it through a stage gate assessment between AGO and FrontierSI at the end of their Proof-of-Concept phase an additional $50,000 and three months will be provided to undertake limited operational testing.

The proposal response should focus on the $100,000 Proof-of-Concept only, operational testing will be discussed later in each successful project.

Intellectual Property

Project Intellectual Property (Project IP) in the capability demonstrators produced during the Project will be owned by the participating partner, in agreement with the lead partner. The AGO is granted a perpetual licence to use any Project IP created for Defence Purposes generally (other than Commercialisation) including internal research, development, education and training. In relation to the use of use of software, source code and project code provided by a project partner, the licence will end at the conclusion of the relevant Project, and the AGO will be required to uninstall and decommission the relevant material promptly following the end of the licence term.

Ownership of Background IP of participating partners for the capability demonstrators will be retained by the participating partners.

AGO will own the project IP in the report required to be produced by participating partners in each demonstrator project.

Technical and Data Support

Many of the challenge topics may be able to use data sources that the successful project partners already have access to. This is the preferred option for projects. Open-source data sets such as SpaceNet may also provide an option – refer to https://spacenet.ai/

Data that have been collected by AGO may be able to be provided to successful partners if required, depending on agreed contractual arrangements.

Additional Resources

A webinar will be completed within two weeks of releasing the call for projects. This webinar will include a presentation of the aims of the AGO Labs program, as well as interactive Q&A with AGO and FrontierSI.

A registration link for the webinar will be provided on the AGO Labs page.

Additional questions can be directed to FrontierSI at agolabs@frontiersi.com.au

Evaluation Criteria

Projects that meet the following criteria will be considered, for further information please refer to the AGO Labs Template found at the AGO Labs website https://frontiersi.com.au/agolabs/

Required

- Project outputs that will address the challenge topic
- Evaluation of the intended approach
- Ability of AGO to access and test outputs iteratively during the project
- Value for money
- Outcomes that can be operationalised beyond the Proof-of-concept project

Constraints and Requirements

- The Australian Government contribution will be limited to $100,000 per project for the Proof-of-Concept
- Projects should be completed in a period of approximately three to six months
- A final project report and presentation is required at the end of the AGO Labs project
- The lead organisation must be a company from Australia or New Zealand

The final project report should outline key findings and recommendations to FrontierSI and AGO. This report will include lessons learned and suggest options for industry partners to engage and work more effectively with AGO. It will also suggest options for AGO to engage and work more effectively with industry through AGO Labs.

Challenge Topic 4: What Vessel is That?

Overview

This challenge aims to investigate whether algorithms can be run across commercial Synthetic Aperture Radar (SAR) satellite data and provide confidence levels for the classification of any given detection for maritime domain awareness.

Problem Statement

The most predominant inhibitor to tipping and cueing for collection using commercial systems is the inability to recognise vessels of interest from SAR imagery.

AGO is interested in a scenario where broad-based tracking systems such as the maritime Automatic Identification System (AIS) provide an overall maritime domain awareness, and then using commercial SAR systems to undertake automated identification of vessel categories. The specific ask is whether it is possible to develop an algorithm that can be run over SAR imagery to provide a classification for vessel categories (with some degree of confidence – at least 60%). Specifically:

- Create an algorithm that will recognise to some level of confidence, a military vessel (at sea) between 50m-200m on commercial SAR broad-area radar with a resolution between 5m-25m.
- Create an algorithm that will recognise to some level of confidence, a fishing vessel (at sea) between 10m-80m on commercial SAR broad-area radar with a resolution of approximately 5m-10m.
- Create an algorithm at that will recognise to a high level of confidence, a military or fishing vessel (in port or harbour) on commercial SAR radar with a resolution between 3m-5m resolution.
- Understand at what resolution detections of different sizes of vessel can no longer be classified to any level of confidence.

Desired outputs: commercial or open-source algorithms that can be hosted on AGO systems to perform SAR analysis.

Data

It is expected that access to commercial SAR maritime datasets will prove challenging. AGO cannot offer any material from its own catalogues. AGO is keen to determine from industry if they have their own SAR assets that could be used or understand the cost implications of procurement of SAR data.

An option would be to use Sentinel SAR data, available from the ESA Copernicus Open Access site at https://scihub.copernicus.eu. It should be noted that these datasets are subject to 'creative commons' terms and conditions that could have implications for an IP developed as part of an AGO funded challenge using this dataset.

Use Cases

The following stories are provided as guidance to challenge responses but should not be considered definitive - AGO is willing to entertain modifications to the following list or additional use cases:

- As an **analyst**, I want to know if the SAR vessel detection I am looking at could be a vessel of interest to me.
- As an **analyst**, I want to know if there is high enough probability that a detection is a vessel of interest to justify a high-resolution imagery collect for further information.
- As a **watch keeper**, I want my operational picture to be updated with the most current and accurate data.
- As a **collector**, I would expect there is justification for collecting additional imagery over a vessel of interest.

Challenge Topic 5: Made You Look

Overview

This topic is a research-focussed investigation around the possibility of encountering (and detecting) malicious modification of commercial imagery product. AGO has pursued synthetic image generation to support model training, but there is a potential dark side to the ability to insert or delete objects from images in ways that are indistinguishable from reality (at least to the Machine Learning system).

How can AGO perform risk mitigation on the collection and dissemination chain to ensure that images have not been modified in uncontrolled aspects of that chain?

Problem Statement

With commercial imagery sources being delivered for high-side analysis via cross-domain systems with low latencies, and a future where AGO could potentially use this to directly feed Machine Learning-based analytic systems, the possibility of malicious modification of imagery in transit is a growing concern.

AGO is continuing to invest in synthetic imagery generation to solve the problem of acquiring training data for ML development; it needs to ensure that this investment is not waylaid through either misdirected or maliciously injected imagery. There are techniques for assurance or integrity which can be applied to parts of the delivery chain AGO controls (such as file hashing) but these can't necessarily be verified for commercial partner systems. For this reason, AGO would like to investigate methods to detect modifications in image pixels or metadata (e.g. changes to the geospatial registration of the image). There are attempts to detect image modification via Benford's Law, but this has not been assessed against satellite imagery.

The desired result of this challenge topic is to provide AGO with an overview of the current state of the art in detecting image modifications, and specifically to assess how well will these work on satellite images with the kind of modifications which are feasible, such as:

- Insertion of fake objects
- Erasure of real objects
- Changes to metadata to reduce the value of any derived data
- Poisoning attacks to images that impact the effectiveness of object detection.

The focus is on detection however and in order to do this AGO will need to know the current techniques available.

Desired Outputs: Report detailing current state of the environment, specifically how to detect disruptions to satellite captured images. A Proof of Concept or demonstration would be useful however AGO is seeking a research effort rather than delivery of an actual capability. Note that AGO has an understanding of best security practices so is not after a security assessment of its supply chain or how to protect assets.

Data

Commercial (Maxar) imagery already obtained for other AGO Labs activities. Will focus on Canberra Airport.

Use Cases

The following stories are provided as guidance to challenge responses but should not be considered definitive - AGO is willing to entertain modifications to the following list or additional use cases:

- As an **analyst** I would like to understand if automated object detection algorithms could be tricked into detecting a non-existent object in a satellite image of Canberra airport.
- As an **analyst** I would like to understand if automated object detection algorithms could be tricked into not detecting a real object in a satellite image of Canberra airport.
- As an **analyst** I would like to understand what technologies or techniques I can use to detect sophisticated addition or removal of objects to commercial satellite images.

Challenge Topic 6: What's Behind that Hill?

Overview

Machine learning algorithms have been developed for detecting and identifying objects in satellite imagery. Object detection models are usually trained using NADIR imagery. This challenge looks to investigate how masking by terrain or structures in off-nadir imagery can be detected.

Problem Statement

In any analysis, the absence of an entity/object is very often, more significant than its presence. However, the presence of objects can be masked by either terrain features or other objects, particularly where the collection geometry is off-nadir. When using satellite imagery to determine the presence or absence of an object there are three potential states:

1. The location of the object is **visible** and the object is **present**
2. The location of the object is **visible** and the object is **not present**
3. The location of the object is **not visible**, therefore the presence of the object is **unknown**.

It is this third state "unknown" that is the subject of interest for this topic.

Using knowledge of terrain (with or without buildings) has the potential to reduce false alarm rates for ML models, but to also contextualise model outputs so that the analyst knows where the model was able to "look" and identify instances where masking has prevented detection of potential objects/entities.

Note that obscuration due to cloud cover is out of scope for this challenge.

Desired Outputs: Commercial or open-source algorithms that can be hosted on AGO systems to perform analysis to determine if a specific location or set of locations will or will not be visible in imagery based on the collection geometry. Ideally this would provide information on partial obscuration as well.

Data

Finding examples of masked terrain suitable for test data will require effort from AGO teams.

Use Cases

The following stories are provided as guidance to challenge responses but should not be considered definitive - AGO is willing to entertain modifications to the following list or additional use cases:

- As an **analyst** I want to know that a single image is able or unable to detect a potential object due to terrain masking by topography, objects or off nadir collection.
  - Example 1 – "X Marks the Spot". For a given location (such as a parking space or aircraft apron) the analyst wants to know if a particular object (such as vehicle or aircraft) is present or not (i.e. aligns to state (1), (2) or (3) above).
  - Example 2 – "Number of Cars in an Area". In order to monitor activity in an area an algorithm is run to count cars for each new image. Nadir images give a sufficiently accurate operational picture. However off-nadir images introduce obscuration of vehicles by buildings, causing the vehicle count to vary based on satellite geometry. So a nadir image may report 100 vehicles; an off-nadir image may report 70 vehicles, which would imply a 30% decrease in activity between the two images. Is it possible to furnish the analyst with a report to advise that there are 70 vehicles visible and 30% of locations previously visible are not visible in this image?
- As an a**nalyst** I want to be able to run an algorithm on a collection of images to determine if a specific geographic point is masked or exposed by terrain or objects.
- As an **analyst** I want to be able to determine the state of a known object (visible or masked) by running an algorithm on a collection of images.

Challenge Topic 7: Detections from Geometry

Overview

Sometimes AGO may know the physical dimensions of an object or facility of interest but might not know its exact geographic location. This can lead to situations where analysts are required to manually search through available imagery, which is time consuming or outright impractical depending on the potential search area.

Problem Statement

Can machine learning techniques identify categories of sites in overhead imagery, based on the input of known dimensions and/or shape of an object or facility? This challenge also looks to understand the computational effort required to scan sets of images across large geographic areas for geometric dimensions.

For such a search algorithm, how specific does the input need to be? Do the search parameters need to have access to precisely measured 3D models? (noting that access to masses of training data is problematic as evidenced by other AGO Labs Challenge Topics) or can input parameters be as simple as:

- Airfields: search for a ground-level (consistent elevation) structure that has a width of 20-30 metres and length of 1000-1500 metres, aligned to any compass point
- Basic Ports: search for a facility neighbouring a water body with coastal/ocean access, that has a pier length of at least 50 metres
- Schools: search for a building with a semi-square footprint, approximately 70m by 70m, with a central courtyard that is approximately 30m by 30m
- Camp: search for an array of tent objects or temporary structures

In a related assessment, do these sorts of searches return a number of 'false hits' (such as a search for runways returning rectangular buildings or sports fields)?

How do collection parameters affect detections? Can we manage skewed shapes caused by imagery collection azimuth or elevation?

Desired Outputs: Commercial or open-source algorithms that can be hosted on AGO systems to perform geometry-based detection.

Data

Ideally, data to be provided by successful challenge participant although AGO may be able to assist with data provision.

Use Cases

- As an **analyst** I want to perform a wide-area search for a facility based on building shape or building dimensions.
- As an **analyst** I want to perform a wide-area search for a facility based on the layout of structures (shapes and dimensions).