

ASSURING AUTONOMY

INTERNATIONAL PROGRAMME

DEMONSTRATOR PROJECT

Final report

ACTIONS: Autonomous Capabilities and Trusted Intelligent Operations in Space

JULY 2022



TECHNICAL REPORT



Project	ACTIONS
Customer	Assuring Autonomy International Programme
Contract	CPL-PRJ-20-01029
Author	Lucy Donnell, Craft Prospect Ltd
Contributors	Hazel Jeffrey, Craft Prospect Ltd Stuart MacCallum, Global Surface Intelligence
Reviewed by	Murray Ireland, Craft Prospect
Approved by	Murray Ireland, Craft Prospect
Rating	Public
Release	1
Date	6 th July 2022



EXECUTIVE SUMMARY

The introduction of onboard autonomy to satellite missions is an ongoing and ever-increasing area of interest across industry and academia. One of the key concerns raised by stakeholders in such missions is the level of trust that can be placed in algorithmic operators versus ground-based human operators. With data analysis and operational decision responsibilities moved upstream and only intermittent ground station contact available to verify these autonomous activities, it is critical that such activities are rigorously assured and can be trusted within some reasonable limits.

The ACTIONS project targeted a demonstration scenario of active fire detection carried out autonomously by an onboard ML component. The driving application was to generate a fire detection alert to emergency response services on the ground, with confidence that the data generated was accurate, truthful, and timely. Data products were also created for downstream commercial applications, supporting the recovery of areas affected by wildfire. Project partners Global Surface Intelligence (GSI) demonstrated a burnt area detection prototype, running inference on data products captured on board to locate fire affected areas suitable for time sensitive recovery activities, such as reseeded with native vegetation.

The following table outlines the objectives met by the ACTIONS project and under which work package they were achieved. The work carried out and results achieved under each work package are described in the sections below, with a final section describing the impact of the project and alignment of future work to the BoK framework.

	Objective	Complete	Justification	Work Package
1	Align other consortium projects to BoK framework where applicable	100%	Work in other commercial and R&D projects has benefited from alignment to BoK guidance (AMLAS, reqs def, SUDA architecture, simulation)	N/A
2	Engage with regulatory bodies across space industry to understand and capture regulatory perspectives	100%	UKSA were engaged early in the project and gave their perspectives on on-board autonomy and the impact on regulations	WP1
3	Implement prototypes of autonomy assurance approaches around disaster application	100%	Prototypes of on-board and service segment autonomous data processing management were created	WP2-3
4	Investigate potential failure modes and test the disaster application in HIL simulation	100%	Hazards and failure modes were identified during requirements capture; autonomy prototypes were tested in real-time HIL simulation	WP2
5	Address and produce guidance for the BoK framework relevant to space	100%	Practical guidance created for 4 BoK objectives	WP4

TABLE OF CONTENTS

EXECUTIVE SUMMARY.....	2
TABLE OF CONTENTS.....	3
1. INTRODUCTION	5
1.1. Document Scope	5
1.2. Project Summary	5
1.2.1. Project Team.....	5
1.3. Version Control.....	5
1.4. Acronyms & Abbreviations.....	6
1.5. Disclaimer	7
2. BACKGROUND	8
2.1. The Challenge	8
2.2. Leveraging Satellite Autonomy	10
3. WORK PACKAGES	12
3.1. WP1: Space Autonomy Assurance	12
3.2. WP2: HIL Assurance Development.....	12
3.2.1. Identification of Hazards	13
3.2.2. Requirements Definition	13
3.2.3. System Safety Requirements.....	13
3.2.4. ML Safety Requirements	14
3.2.5. ML Component Development	15
3.2.6. ML Component Testing and Verification.....	15
3.2.7. HIL Simulation Testing	16
3.2.8. Evaluation of HIL Simulation Testing Results	19
3.2.9. Downstream Application.....	21
3.3. WP3: MBSE Process Assurance	24
3.4. WP4: Final Evaluation and Guidance.....	26
3.4.1. Defining Operating Scenarios (1.1.3).....	26
3.4.2. Implementing Requirements Using ML (2.3)	26

3.4.3.	Using Simulation (2.7)	26
3.4.4.	Defining Understanding Requirements (2.2.1.2)	26
4.	IMPACT	27
4.1.	BoK Alignment	27
4.2.	Engagements	27

1. INTRODUCTION

1.1. Document Scope

This document provides background information about the emergency of wildfire events, and the leveraging of satellite autonomy to improve outcomes by enabling a faster emergency response. The work carried out to meet the objectives of the ACTIONS demonstrator project is described and the key deliverables and impact of the project are also outlined.

1.2. Project Summary

ACTIONS gives insight into autonomy assurance in small space systems. To do this, we use our existing demonstration systems and processes developed for implementing autonomy on-board satellites. We deliver a demonstrator including elements of MBSE, on-board flight software, simulation, and hardware-in-the-loop (HIL) testing for specific machine learning (ML) algorithms and other components. We use a disaster response scenario as the driving application. This allows us to focus on assurance of a use case with specific time and quality constraints. Given the limited resources of small satellites and the sparse opportunities for data capture, autonomy offers significant improvements in utilisation and timeliness of service to end users.

This work complements and aligns with projects and products from Craft Prospect and partners, leading to an in-orbit demonstration in 2023. We intend to align partner projects to the Body of Knowledge framework, in addition to existing space standards and industry best practice. Previous stakeholder engagement across the space industry is leveraged to understand current and future needs in autonomy. This ensures space is considered within the AAIP BoK, informs further applied work in spacecraft autonomy, and potentially integrates with frameworks for future missions.

1.2.1. Project Team

Lucy Donnell, Murray Ireland and Hazel Jeffrey

Craft Prospect Ltd, Glasgow, UK


Richard Hawkins and Chiara Picardi

University of York, York, UK

Stuart MacCallum, Mark Howie and Freddie Hunter

Global Surface Intelligence, Edinburgh, UK

1.3. Version Control

Release	Author	Date	Change Control	Approved
1	Lucy Donnell	July 2022	Public release version	

1.4. Acronyms & Abbreviations

AAIP	Assuring Autonomy International Programme
ADCS	attitude determination and control system
AMLAS	Assurance of Machine Learning in Autonomous Systems
ARCADIA	architecture analysis & design integrated approach
BoK	Body of Knowledge
CPL	Craft Prospect Ltd
ESA	European Space Agency
EO	Earth observation
ER	emergency response
FPS	frames per second
GSI	Global Surface Intelligence
HIL	hardware in loop
ITA	Institute for Trusted Autonomy
MBSE	model based systems engineering
ML	machine learning
PPS	pixels per second
QPS	queries per second
ROI	region of interest
REQ	requirement
SIL	system in loop
SUDA	sense, understand, decide, act
UKSA	UK Space Agency
UoY	University of York
V&V	verification and validation

1.5. Disclaimer

Craft Prospect Ltd. does not provide any warranty whatsoever, whether expressed, implied, or statutory, including, but not limited to, any warranty of merchantability or fitness for a particular purpose or any warranty that the contents of the item are error-free. In no respect shall Craft Prospect Ltd. incur any liability for any damages, including, but not limited to, direct, indirect, special, or consequential damages arising out of, resulting from, or in any way connected to the use of this document, whether or not based upon warranty, business agreement, tort, or otherwise; whether or not injury was sustained by persons or property or otherwise; and whether or not loss was sustained from, or arose out of, the results of, the item, or any services that may be provided by Craft Prospect Ltd.

2. BACKGROUND

2.1. The Challenge

Wildfires are an unfortunately common and often catastrophic occurrence in many parts of the world. In the USA, wildfires remain an ongoing concern in several states (Table 1). 2020 and 2021 were the worst years for wildfires in the USA in at least 10 years (Figure 1 and Figure 2). In Oregon in 2020 alone, more than 400,000 hectares of land were burned, thousands of homes were destroyed and 11 lives lost¹.

Table 1 – States at high extreme wildfire risk, 2021².

Rank	State	Est num properties at risk
1	California	2,040,600
2	Texas	717,800
3	Colorado	373,900
4	Arizona	242,200
5	Idaho	175,000
6	Washington	155,500
7	Oklahoma	153,400
8	Oregon	147,500
9	Montana	137,800
10	Utah	136,000
11	New Mexico	131,600
12	Nevada	67,100
13	Wyoming	36,800

¹ Newburger, Emma. ‘At Least 33 Dead as Wildfires Scorch Millions of Acres across Western U.S. — “It Is Apocalyptic”’. CNBC, 12 September 2020. <https://www.cnbc.com/2020/09/12/fires-in-oregon-california-and-washington-spread-death-toll-rises.html>.

² ‘Facts + Statistics: Wildfires | III’. Accessed 21 April 2022. <https://www.iii.org/fact-statistic/facts-statistics-wildfires>.

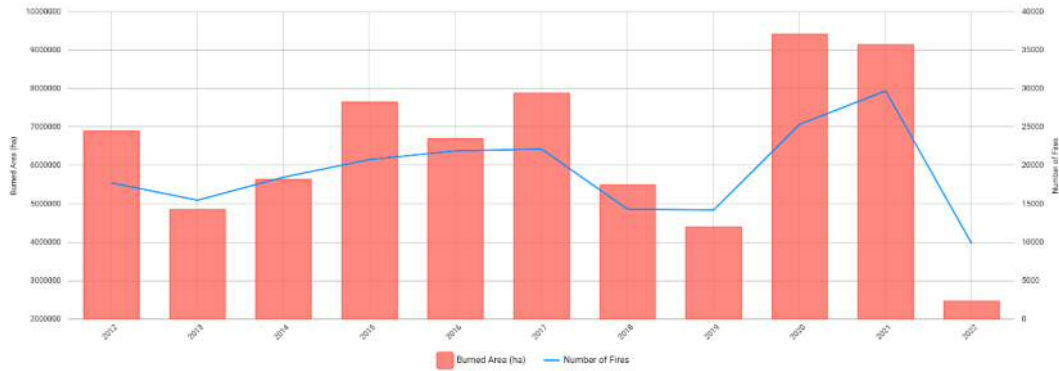


Figure 1 – Estimated burned area and number of fires in the USA, 2012 to 2021³.




Figure 2 – Annual number of acres burned in wildland fires, 1980-2021⁴.

The need to respond quickly and efficiently to wildfires is well established. Services such as the Fire Information for Resource Management System (FIRMS)⁵, the Global Wildfire Information System (GWIS)⁶ and the Copernicus Emergency Management System (EMS)⁷ have been created to provide early warnings, statistical data and coverage maps for wildfires. This allows the response to active wildfires to occur as quickly as possible upon detection. It also allows the occurrence of wildfire to be predicted and mitigated where possible. Data on burnt areas is typically also provided to enable damage assessment and recovery planning.

Such services rely heavily on satellite data to provide the perspective, spectral content and temporal frequency needed for regular and accurate detection and reporting of wildfires and

³ 'GWIS - Statistics Portal'. Accessed 21 April 2022.

<https://gwis.jrc.ec.europa.eu/apps/gwis.statistics/estimates>.

⁴ 'Facts + Statistics: Wildfires | III'. Accessed 21 April 2022. <https://www.iii.org/fact-statistic/facts-statistics-wildfires>.

⁵ NASA. 'FIRMS - Fire Information for Resource Management System'.

<https://modaps.eosdis.nasa.gov/map/>.

⁶ 'GWIS - Welcome to GWIS'. Accessed 21 April 2022. <https://gwis.jrc.ec.europa.eu/>.

⁷ Copernicus. 'Copernicus Emergency Management Service'. <https://emergency.copernicus.eu/>.

burnt areas. As these services rely on existing missions, however, they are subject to the limitations of these missions in terms of visit frequency, information latency and quality of data. For example, FIRMS reports a lead time of 3 hours from observation (not the fire actually starting or being observable) to distribution on ground⁸, geolocation precision of 375 m⁹ or 1 km¹⁰ and a commission error of 1.2%¹¹. The source missions for FIRMS (Terra, Aqua, Suomi NPP and NOAA-20) have a revisit time of between 14 hours and 2 days. This makes the worst-case scenario for a detection response time around 51 hours, assuming a fire becomes observable immediately following a satellite pass.

Wildfires can spread as quickly as 10.8 kph in forest and 22 kph in grassland¹². The damage to infrastructure and loss of life which could potentially be caused in the time it takes to report an active wildfire is huge. While emergency services do not rely exclusively on platforms such as FIRMS or Copernicus EMS, the ability to provide warnings even a few hours earlier could make a huge difference to the preservation of human, animal and plant life and infrastructure.

2.2. Leveraging Satellite Autonomy

The lead time on fire alerts can obviously be reduced by increasing the revisit frequency of any observation satellites or deploying a constellation intentionally sized and designed to meet specific revisit and latency requirements. However, there is still the need to process significant volumes of observation data on the ground, identify the presence, location and other salient details of wildfires, and disseminate this information to end users such as emergency services. Shifting this processing upstream and providing observing satellites with the ability to identify fires in observation data at the edge enables:

Rapid tagging and filtering of data: prioritising data which is believed with a high degree of confidence to include wildfires.

Alert generation: extracting the salient information from the raw observation data, such as the detection time, location and size of identified wildfires. Such alerts can then be prioritised at the front of the downlink queue or transmitted to end users via more frequent, lower-bandwidth ground station passes.

⁸ NASA. 'FIRMS - Fire Information for Resource Management System', n.d. <https://modaps.eosdis.nasa.gov/map/>.

⁹ 'VIIRS I-Band 375 m Active Fire Data | Earthdata'. Accessed 21 April 2022. <https://earthdata.nasa.gov/earth-observation-data/near-real-time/firms/viirs-i-band-active-fire-data/>.

¹⁰ 'MCD14DL | Earthdata'. Accessed 21 April 2022. <https://earthdata.nasa.gov/earth-observation-data/near-real-time/firms/mcd14dl/>.

¹¹ Schroeder, Wilfrid, Patricia Oliva, Louis Giglio, and Ivan A. Csizsar. 'The New VIIRS 375 m Active Fire Detection Data Product: Algorithm Description and Initial Assessment'. *Remote Sensing of Environment* 143 (March 2014): 85–96. <https://doi.org/10.1016/j.rse.2013.12.008>.

¹² Billing, P. 'Otways Fire No 22 - 1982/83 - Aspects of Fire Behaviour'. Fire Research Branch. Victoria Department of Sustainability and Environment, June 1983.

Verification data generation: creating ancillary data products such as image thumbnails, detection reports and augmented visualisations. These can be packaged with alerts for verification or queued closely behind alerts during downlink.

Data reduction: without knowing which parts of raw data are valuable and which aren't, raw data can only be losslessly compressed for downlink, retaining a large memory footprint and using up valuable downlink bandwidth. If the valuable regions of data can be identified, these can be retained at full quality while other regions are discarded, reduced or lossy compressed.

Responsive tasking: detected wildfires can be revisited on subsequent passes by autonomously tasking the satellite to target and re-acquire locations of detections. This enables ongoing monitoring of fires and change detection to determine rates of spread.

This on-board processing is enabled using machine learning algorithms trained to recognise fires in observation data and report the locations of these fires and the confidence of the prediction. With trust in these reports, early warnings can be generated on-board and delivered to the ground and emergency services ahead of large volumes of raw data, whose content and therefore value is otherwise unknown.

3. WORK PACKAGES

This section describes the work packages that were undertaken to deliver the outcomes of the ACTIONS demonstrator project.

3.1. WPI: Space Autonomy Assurance

The autonomy in the ACTIONS mission is non-mission-critical: decision-making is open-loop and does not affect the behaviour of systems such as the payload, ADCS and comms (excluding the content of the downlink queue). The autonomous component affects the number and content of the data products that are created on-board and the makeup of the downlink queue, populated by these products. As such, it is safety critical in the context of the larger space system, as the end use of the data products – emergency response – is directly related to safety of life. Figure 3 illustrates the flow of data through the ACTIONS fire alert system.

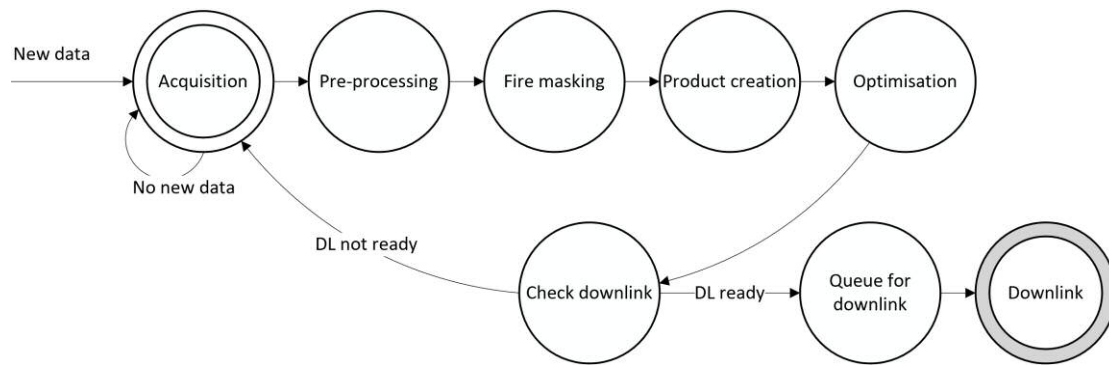


Figure 3 – ACTIONS processing chain dataflow diagram.

In the first stage of the project, workshops were held with industry contacts. AAC Clyde Space were engaged for their views as a manufacturer of satellite platforms and subsystems, with future interest in autonomous constellation operations. Bright Ascension were engaged for their views as a developer of satellite flight software, with ongoing activities in development of autonomous software and assurance of spacecraft software.

To gain insight into the impact on regulation and licencing of satellite autonomy, the UK Space Agency was also engaged. Their view was that such autonomy has minimal impact on the licencing process as it currently exists. The primary impact would manifest in the evidence required to show that the satellite operations are safe and would not harm another spacecraft in orbit.

3.2. WP2: HIL Assurance Development

Throughout this work package, the following goals were achieved:

- Identify the key hazards to life brought on by utilising autonomy in a wildfire detection mission.
- Define requirements which must be met to deliver a safety assured end-to-end early warning system for wildfire events.

- Understand the role of machine learning (ML) in the autonomous mission and identify the requirements of ML components such that system-level requirements are met.
- Deliver a demonstration system for autonomous wildfire detection and reporting, tested in a realistic mission simulator.
- Develop and test a commercial application of the ACTIONS mission, where data products generated onboard are used for ground based burnt area detection to support the recovery of wildfire affected areas.

The following sections outline the work carried out during the HIL Assurance Development: Identification of Hazards, Requirements Definition, ML Component Development, ML Component Testing and Verification, Simulation Testing, and the Evaluation of Simulation Test Results. Following this, the development of the commercial application (burnt area detection) is described.

3.2.1. Identification of Hazards

There is a need to understand the hazards and failure modes of the space system, to identify the causes by which the mission will fail. Failure in the ACTIONS mission is defined as the failure to respond to a fire in sufficient time to contain it and prevent loss of life.

The autonomy hazards of the ACTIONS mission were identified as the following:

- A wildfire is reported at a location where there is none, leading emergency services to waste time and resources responding to it. Other, real, fires may be left unchecked as a result. This can be caused either by a false positive or incorrect geolocation.
- No wildfire is detected at a location where a wildfire does exist. The wildfire is then left unchecked and uncontrolled. This can be caused either by a false negative or incorrect geolocation.
- A wildfire and its location are accurately reported, but too late for it to be contained.

These hazards form the basis of the safety requirements for the mission. To consider the mission outputs “safe” within some bounds, requirements which address these hazards must be defined and met. These safety requirements can be functional or performance requirements; the key consideration is that they are realistic and verifiable.

3.2.2. Requirements Definition

The system safety requirements defined for the ACTIONS system necessitate that the fire alerts generated in-orbit and sent to the emergency services on the ground are accurate, truthful, and timely.

3.2.3. System Safety Requirements

Missed detections, or misdirection of emergency services to attend non-fires both pose a risk to property, the natural environment, and potentially to human life. Four system safety requirements were defined in response:

- REQ-SAFE-ER-1 - The Emergency Response Service shall determine the location of a visible active fire within 200 m of its true location.

- REQ-SAFE-ER-2 - The Emergency Response Service shall inform emergency services of a visible active fire within 3 hours of it starting.
- REQ-SAFE-ER-3 - The Emergency Response Service shall positively identify 95% of all visible active fires acquired by the satellite instrument within the area of interest.
- REQ-SAFE-ER-4 - The Emergency Response Service shall falsely indicate visible active fires in the area of interest at a rate not exceeding current fire alert service.

3.2.4. ML Safety Requirements

The defined system safety requirements were then allocated and interpreted for the ML component specifically. Understanding the make up of the ML component and its interfaces within the system is key to this allocation. Figure 4 visualises where the ML component interfaces exist within the ACTIONS emergency response system.




Figure 4 – System interfaces of ML component.

The neural network model developed for the ML component performs semantic segmentation, carrying out fire detection at the pixel level. The model outputs masks, where each pixel (representing a specific area on the ground) is labelled as fire or non-fire.

Understanding what the ML model development data contains was crucial in understanding what the output detections represent. Important features of the data used to train the ACTIONS neural network model were:

- Truth labelling: The process by which the training data has been labelled as containing fire or non-fire pixels determines how detection will be carried out by the model.
- Data format: The metre per pixel resolution of the imagery informs the location accuracy of the detection mask.

The ML component safety requirements for the ACTIONS system were defined as follows:

- REQ-SAFE-ML-1 - All points of the mask generated by the ML component shall be less than 6 pixels outside the boundary of the area of the real fire.
- REQ-SAFE-ML-2 - The ML component shall correctly identify the presence of a fire that satisfies the Schroeder¹³ conditions in a frame for 95% of real fires.

¹³ The sensor tuned conditions for active fire detection set out by Wilfrid Schroeder, Patricia Oliva, Louis Giglio, Brad Quayle, Eckehard Lorenz, and Fabiano Morelli. Active fire detection using Landsat-8/OLI data. Remote Sensing of Environment (Elsevier), 185:210 – 220, 2016. ISSN 0034-4257. doi:10.1016/j.rse.2015.08.032.

- REQ-SAFE-ML-3 - The ML component shall not identify the presence of a fire in a frame where there is not a real active fire more than 52¹⁴ times per month.
- REQ-SAFE-ML-4 - ML performance requirements shall be satisfied for all data across the range of featured present in the operating scenarios.

Specific details of ML requirements were iterated over as the development data and architecture of the ML model were each defined, also considering the operational scenarios of the system. These requirements were then implemented throughout the development lifecycle of the ML component (data management, model development, model verification and model deployment).

3.2.5. ML Component Development

When selecting a model for deployment onboard a small satellite, domain specific factors inform model architecture selection. Due to memory and power constraints, a smaller model architecture was necessary for the ACTIONS mission. Inference speed must reach a certain threshold for successful on-board processing of sensor data. Selecting a simpler model, with fewer parameters, also serves to make fault diagnosis more straightforward.

Assuring the quality of the alerts generated by the ML component of the ACTIONS system is critical. Early warnings are no use to emergency services if the information they contain is incorrect or inaccurate.

An ML model is often described as a black box due to the extremely complex mathematical equations that determine the inference process. Clear documentation of model development steps, justification for choices made and meaningful evaluation of performance are all crucial for making the model explainable and instilling trust in it. The various assurance artefacts generated when following the AMLAS process throughout development of the ACTIONS ML component are valuable for communication to customers and partners and building trust in the ML component which enables the autonomy of the emergency response system.

3.2.6. ML Component Testing and Verification

During the testing and verification stage, the model was tested against various datasets. Each dataset generated was documented in the AMLAS artefact: Data Generation Log, which also justified the satisfaction of the key data requirements of relevance, balance, completeness and accuracy.

The results of executing the ML model using labelled test data, and the sufficiency of these results in meeting the ML safety requirements, were documented in the AMLAS Artefact Internal Test Results. At this stage, the model was also tested on unlabelled data. This was to assess performance on continuous data from the ROI, which was not available within the labelled development set. Testing was an iterative process and model development steps were revisited to optimise model performance on the test data.

¹⁴ NASA FIRMS was considered as the gold standard for FPs (at 52 instances a month), therefore equivalent or better performance is safe.

Table 2 provides the performance metrics for the ML component generating fire masks for the labelled test dataset.

Table 2 – Performance metrics for fire masking component against labelled test dataset.

Metric ¹⁵	Value	Notes
Model accuracy – MeanIoU	93%	Very good result for semantic segmentation (indicating that the model output masks are very similar to the label masks).
Model accuracy – true positive	100%	Excellent result (for inference on labelled internal test set)
Model accuracy – true negative	99.2%	

After achieving sufficient results, further verification steps were taken. Verification data was generated outside of the model development and testing process. This is a feature of the AMLAS framework which helps to assure performance by pushing the model to expose limitations and broadening the test data to overcome any potential blind spots of the developer.

3.2.7. HIL Simulation Testing

The ML component was then deployed in a simulated environment with target hardware in the loop across a set of defined operational scenarios.

Hardware-in-loop (HIL) or system-in-loop (SIL) simulation involves the use of physical hardware interfaced with a real-time software simulation. Simulation is required due to the challenges in both feasibility and cost of testing space systems in a realistic operating environment. Instead, simulations can be used to provide realistic orbital, attitude and electromechanical models which can feed flight telemetry and ephemeris data to hardware and software components which will ultimately be deployed in space.

Simulation tools and techniques used in ACTIONS include:

- Mission-level analysis and planning tools to determine constellation size and instrument footprint in order to meet specific revisit requirements – combined with the on-board autonomy, this allows the response time safety requirement [REQ-SAFE-ER-2] to be met.
- Testing of the ACTIONS early warning processing system in real-time, to ensure data throughput is sufficient to meet the requirement for real-time processing.
- Use of a sensing environment model to feed realistic optical data to the simulated multispectral instrument.
- Feeding of realistic ephemeris data and sensor telemetry to the processing system, allowing creation of location-accurate georeferenced products and reporting of fire locations in alerts.

¹⁵ Note that true positives and true negatives allow for a 6-pixel margin of error to account for 200 m geolocation accuracy

An orbit of a single satellite over Oregon (the ROI), was simulated on a day when multiple fires occurred. Figure 5 illustrates the data captured during the simulation.




Figure 5 – Simulated single pass over Oregon with several active fires captured.


Table 3 – Performance metrics for ACTIONS fire masking component in simulation.

Metric	Value	Notes
Model throughput (FPS)	0.3	Sufficient to meet real-time processing requirements
Model throughput (QPS)	413.5	
Model throughput (PPS)	952,680	
Fire masking component latency	3.57 s	
Information latency	4.84 s	


The ML component was deployed in the simulation and the emergency response data products were created and optimised in real-time, validating latency and throughput requirements. Table 3 provides the latency and throughput metrics for the ML component during simulation testing. The end-to-end processing time of an instrument frame (raw data to stored data products) is sufficient to keep pace with the equivalent framerate of the instrument payload (5 seconds).

An example of the data products generated by the ACTIONS system is illustrated in Figure 6. Level-0 and Level-1 multispectral images provide full, unbiased data products to enable ground-based V&V and re-training, as well as a variety of secondary ground applications. Georeferenced pixel fire masks extracted from L1 products can enable precise geolocation of fires – down to 30m. Level-4 products provide extremely lightweight text alerts – containing only the salient details – supported by low-resolution annotated thumbnails for additional assurance. Emergency services can then:

1. Act as soon as the alert is received.
2. Quickly receive the L4 thumbnail for visual validation of the alert.
3. While on-route, obtain more precise details on fire location through the L3 mask.
4. Later receive the full-size source image data for V&V and quality control.




a) Level-1 false colour image.



b) Level-3 fire mask.

FIRES DETECTED
LARGE fire:
Time 20200908-065847.184
Location (123.4094°S, 41.8675°E)
Area 40.08 sq km

c) Level-4 alert message.




d) Level-4 verification thumbnail.


Figure 6 – Sample data products created by ACTIONS system.

3.2.8. Evaluation of HIL Simulation Testing Results

The ability of emergency services on the ground to respond quickly hinges on the L4 alerts being downlinked quickly. In existing solutions such as FIRMS and Copernicus EMS, L0 data must be downlinked in order of acquisition and processed on the ground to extract the salient fire information. In the ACTIONS system, the L4 alerts are prioritised and then followed by the ancillary products to provide assurance and V&V. This is enabled by the ability to autonomously label all data products derived from a discrete payload acquisition. This prioritisation is illustrated in Figure 7.



a) Traditional downlink queue for payload data.



b) Intelligent downlink queue for payload data and derivative products.

Figure 7 - Comparison of traditional downlink queue for payload data with intelligent queuing. Products indicating the presence of wildfires are marked in red, while those without are blue.

The benefits to data latency can be quantified. Consider a 10 Mbit downlink and 50% probability of fire being present in a captured image frame. Representative file sizes for each product type after compression are given in Table 4. In an 8-minute ground station pass, assuming optimal conditions and minimal connection overheads, 30 L0 images can be downlinked in a traditional downlink scenario. This has two major issues:

Assuming all new on-board data is downlinked and neglecting the timeliness of the acquisition operations, images showing wildfires could have a downlink latency of up to 8 minutes.

The assumption that all new on-board data is downlinked may be incorrect, and more recent data may need to wait for a subsequent ground station pass before downlink.

Table 4 – File sizes for each product type after compression.

Product	Average file size
L0 multispectral image	20.4 MB
L3m pixel mask	23 kB
L4a alert	5 kB
4m thumbnail image	50 kB

With autonomous on-board queue management, the lightweight alerts are prioritised and the bulky L0 source data is moved to the back of the downlink queue. Only L0 and L3m products showing fires are downlinked in this scenario (50% of all frames). Taking the baseline downlink’s target of 30 L0 products, the ACTIONS system can downlink all fire alerts in 0.12s and all fire alerts, verification thumbnails and geolocation masks in 1.56s. The remainder of the downlink bandwidth can be used to retrieve L0 or L1 data products for V&V and retraining. These results are illustrated in Figure 8. While a much larger number of files has been created and downlinked by the ACTIONS system, these files are smaller in footprint and add enormous value to the end users of the mission.

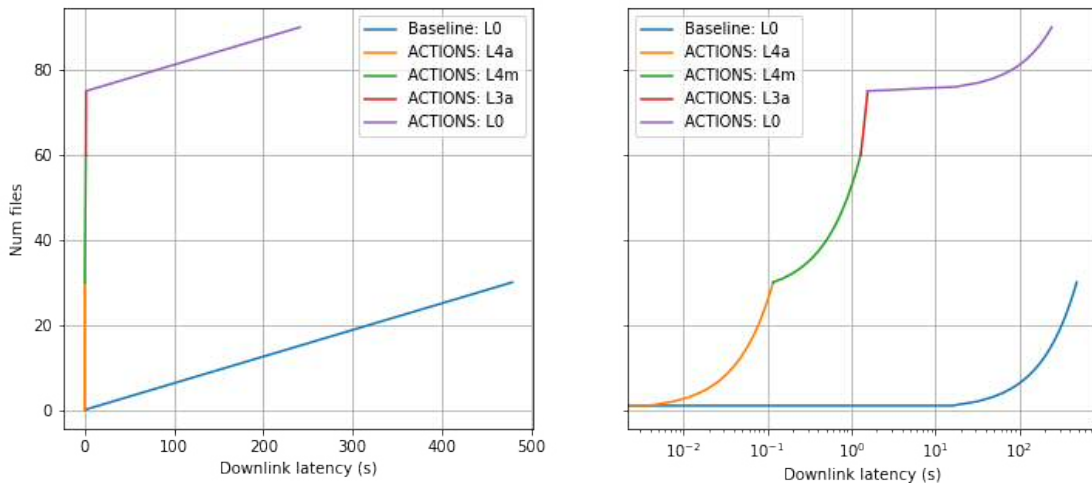


Figure 8 – Comparison of downlink latency for each product file, for baseline and optimised ACTIONS queues. Shown in linear (left) and logarithmic axes (right).

The performance of the model concerning compliance with safety requirements was evaluated by eye. This was due to the absence of available labelled data suitable for simulation testing. This evaluation process was carried out by visualising various overlays of the test data and model masks, enlarging areas of the image for close examination.

During evaluation, the model was considered to have made a mixture of correct and erroneous predictions whenever the case appeared ambiguous, based on factors such as context and intensity of the specific pixels classified as fire or non-fire. Alternative approaches were taken so that results could be recorded implementing a low, moderate and high threshold for determining whether the model was correct when a classification appeared ambiguous and was impossible to verify.

Compliance with the mission safety requirements is addressed in Table 5. Three requirements are fully complied with and the final requirement on false positives partially complied with, the uncertainty indicating a need for further analysis of truth data. Compliance with these requirements indicates that the mission is safe in the context of its early warning application goals.

Table 5 – Compliance matrix for mission safety requirements.

Requirement	Compliant	Evidence
The Emergency Response Service shall determine the location of a visible active fire within 200 m of its true location.	Yes	30 m resolution available in fire mask, geolocation accuracy sub-50 m for test
The Emergency Response Service shall inform emergency services of a visible active fire with 3 hours of it starting	Yes	Requirement met with 188 satellites and single ground station, using intelligent downlink queue. Smaller constellation size possible if aiming to match FIRMS latency only (12 hours).
The Emergency Response Service shall positively identify 95% of all visible active fires acquired by the satellite instrument within the area of interest	Yes	False negatives calculated at 0.76%, yielding 98.24% true positives
The Emergency Response Service shall falsely indicate visible active fires in the area of interest as less than 52 instances per month	Partial	Depending on threshold in validation approach, false positives in simulation tests are either 53 (moderate threshold) or zero (low threshold)

3.2.9. Downstream Application

The ACTIONS demonstrator mission features a ground segment where data products were created for a downstream commercial application, supporting the recovery of areas affected by wildfire. Project partners GSI demonstrated a burnt area detection prototype, applied to data products captured onboard to locate fire affected areas suitable for time sensitive recovery activities, such as reseeded with native vegetation.

This commercial application of the ACTIONS mission does not have safety-critical requirements, i.e., there is no direct threat to human life from inaccuracies in the output products. As a ground segment, it is also inherently more easily adaptable post-launch, than the on-board system. Combined, these two features reduce the levels of assurance required in the ground segment, relative to the ACTIONS emergency response system.

The following sections describe the development process of the burnt area detection application of the ACTIONS mission.

3.2.9.1. ML Component Assurance

Although not safety-critical, all commercial applications require levels of accuracy and assurance that enable robust products of value to be delivered to the end-users. Without some level of assurance, the commercial value of downstream products will be reduced and may ultimately damage the financial health of both provider and end-user.

Implementation of the AMLAS framework is expected to lead to a higher quality end product but its rigorous implementation also bears a financial cost to the provider, without necessarily creating a higher financial value product. Consequently, levels of assurance of commercial applications can be considered as a trade-off between rigour and cost and will vary from application to application.

The AMLAS framework has been considered throughout the development of the ACTIONS commercial application. Each stage of the AMLAS framework was considered to varying degrees, with main areas of focus on: Data Management, Model Verification, and Model Deployment.

In terms of ML Safety Assurance and ML Requirements Assurance, discussions with potential end-users were conducted to establish the initial commercial assurance requirements from which the ML requirements could be derived, based on the overall system design.

The AMLAS framework incorporates an iterative approach, with feedback resulting in earlier stages being revisited as development proceeds. This resulted in changes to requirements to improve clarity and testability based on better understanding of the system and user-requirements.

3.2.9.2. ML Component Development: Data Management

Data management focused on identifying and preparing a suitable reference data for training and verification, with assurance addressing relevance, balance, completeness and accuracy.

Key elements included:

- Use of satellite imagery equivalent to the onboard system design.
- Coverage of the relevant geographic area and land cover types.
- Coverage of a range of fire events (severity, seasons, wildfire/prescribed, forest types, etc.).
- Use of a well-established and understood dataset.
- Balance of a range of burn severity levels and forest types, to provide assurance under different conditions.

Observed limitations (e.g., time-lag to post-fire images, bias towards clear-sky) in the resulting datasets, were acknowledged and potential routes to improvement or mitigation were identified.

3.2.9.3. ML Component Verification

Model verification used this dataset to establish that the ML requirements were met. This verification stage did not fully comply with the AMLAS framework, as the verification dataset was not truly independent, in that it was developed alongside the training dataset by the same team. Although not truly independent, best efforts were made to maximise independence (e.g., sampling from distinct fire events to minimise effects of spatial correlation), with decisions on how to split the available reference data made prior to model development commencing.

Further qualitative assurance was provided through the full system demonstration of the evolution of a single large fire event over time. A second level of assurance such as this, can serve as useful and potentially more efficient method of testing the impact (if any) of limitations in the training and verification data, without incurring the full expense of augmentation/redesign. Results from secondary qualitative assurance may help guide and streamline further iterations of model development within the AMLAS framework.

3.2.9.4. ML Component Deployment

A third level of assurance was provided through integration testing, where the focus was testing under more extreme but also realistic active fire conditions. Issues of false positive detection of damage under conditions of thick smoke/cloud were raised through these additional qualitative assurance stages. These highlight the need for a further iteration through the stages of AMLAS framework, potentially revisiting the Data Management and Model Verification stages to augment or redesign the training data, or potentially considering changes to the wider system (e.g., introducing an independent smoke/cloud detection stage).

Integration testing also provided additional assurance of the onboard emergency response system, allowing verification of the following: onboard geolocation, the locations of active fire detections, and that outputs from the onboard system fit within the ground segment system design (e.g., the correct satellite channels were present).

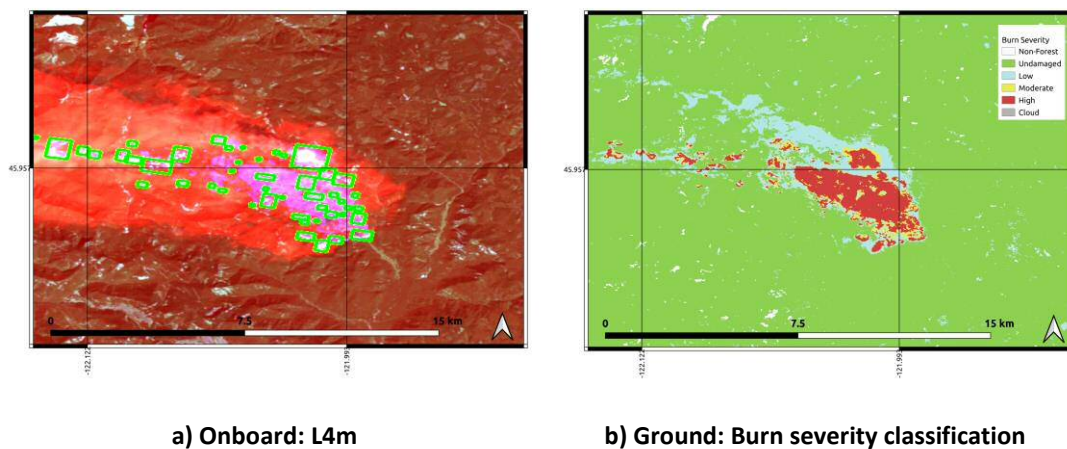


Figure 9 shows a comparison of a fire detection output generated from the onboard segment and a burn severity output generated from the ground segment of the ACTIONS mission.

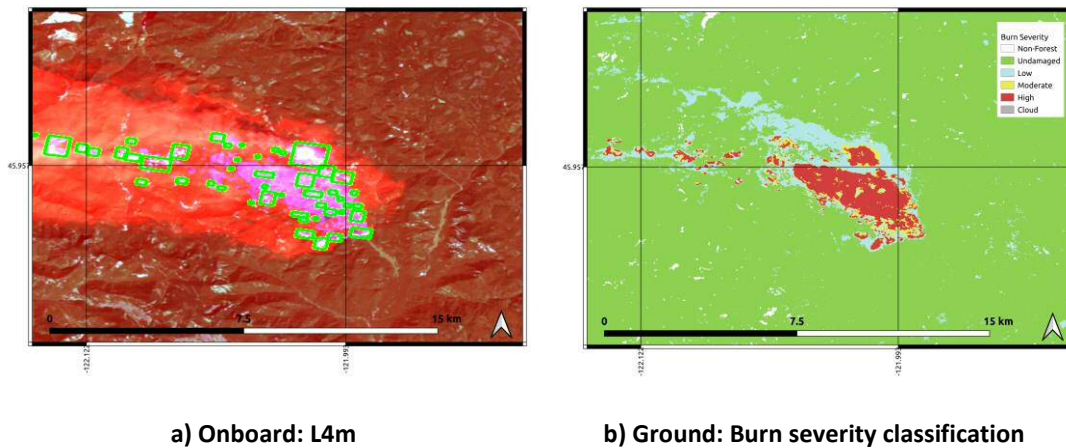


Figure 9 – Side-by-side comparison of outputs from onboard and ground segments.

3.3. WP3: MBSE Process Assurance

Model-based systems engineering (MBSE) was applied to the ACTIONS demonstration project. It spanned the space segment (the emergency response service system) and the ground segment (where data products generated on-board were downlinked and used in burnt area detection for commercial applications supporting wildfire recovery).

MBSE is a methodology that uses one large integrated visual model to develop, document and communicate the requirements and behaviour of a complex system. The traditional systems engineering method is very document-driven and manual, where engineers must manually create entirely separate documents capturing information such as requirements, behavioural architecture and failure mode analysis. Each of these documents must be reviewed and maintained separately, however, it is critical to ensure that information traceability is extremely accurate. If not, there is a high potential for missing a mission-critical failure mode or requirement in the design flow-down and implementation phase.

Following the MBSE process allows the mission developer to visualise and communicate all of the system entities and actors and the physical and functional interactions between these for multiple different use cases. This starts as a definition of the very high-level mission functionalities and is linked down to the physical architecture level.

The benefit of this link is that it provides consistency and transparency when defining more detailed aspects of the mission or system design. When an element is changed at the higher level the user alerted to the impact of the change within the rest of the system which, if accepted, is then automatically implemented within the various layers of the model. In theory this allows the user to identify potential hazards and/or issues more easily and earlier in the design process, leading to a more assured and cost-effective design. Additionally, within the various MBSE tools available the user is required to define a functional or physical link between the system components or activities. This both ensures that no unnecessary elements are created and highlights any potential hazards or failures likely to occur between system mode transitions or with data input/output to a particular system activity.

For the ACTIONS project the MBSE tool of choice was Capella. Following the ARCADIA methodology implemented within Capella, the following MBSE approach was utilised:

- Operational Analysis: Define Stakeholder Needs, Environment, involved entities and actors.
- System Analysis: Identify the system boundaries and define System Functions
- Logical Architecture: Define the functional flow of the system

Within the operational and system analysis stage, swim-lane diagrams were utilised to capture the system behaviour. This was then flowed down to the logical architecture level which focused on modelling the dataflow through the system and identifying the failure modes associated with the functional flow of the system. This was an iterative process, with changes identified to the operational and system analysis information following logical analysis of the system. The relevant information and views from the MBSE model were exported to HTML at various points in the iterative cycle which were reviewed with project stakeholders. These reviews identified potential hazards within the ML component implementation and informed the final design. Conversely, the reviews identified inaccurate modelling within the MBSE model and ensured that all project stakeholders had a unified understanding of the system.

The final output of the ACTIONS MBSE model and the general process followed is depicted in Figure 10.

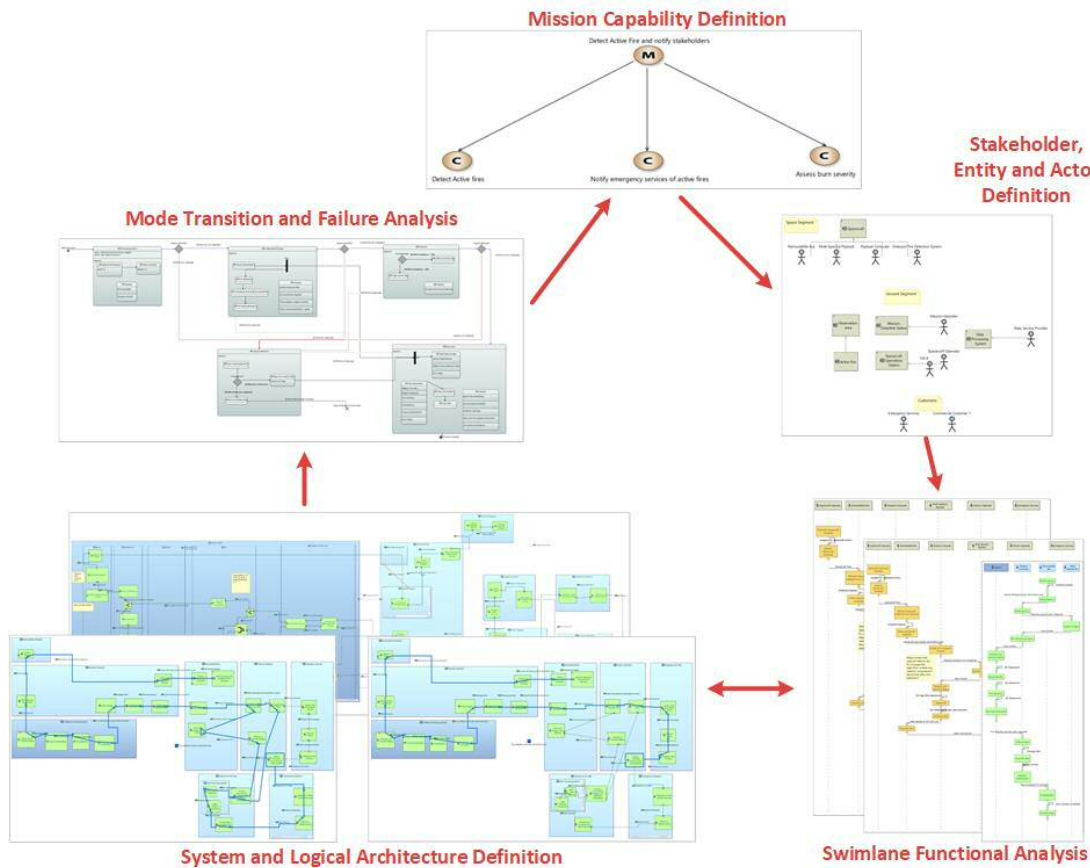


Figure 10 – ACTIONS model-based systems engineering process and outputs.

3.4. WP4: Final Evaluation and Guidance

Contributions were made to the AAIP Body of Knowledge (BoK) to provide practical guidance relevant to the space and earth observation domains, and beyond.

Other project outputs include:

- On-board Autonomy Assurance Prototype Report (emergency response application)
- Service Segment Assurance Prototype Report (commercial application)
- MBSE Process Assurance Report
- A document containing the full collection of assurance artefacts created during the ACTIONS emergency service response application demonstration. This details how the AMLAS process was followed from beginning to end.
- A white paper, with contributions from all project partners, titled 'Autonomy Assurance for Small Space Systems'.

Contributions were produced for the following sections of the AAIP BoK, outlined below.

3.4.1. Defining Operating Scenarios (1.1.3)

Practical guidance was contributed for defining operating scenarios of a small space system, typically composed of space, ground and service/user segments. The process of defining operational scenarios for the ACTIONS emergency response service is described.

3.4.2. Implementing Requirements Using ML (2.3)

Practical guidance was contributed for Implementing Requirements Using ML. The development of the ML component for the ACTIONS emergency response service was used as an example of defining ML component specific safety requirements and following the AMLAS process to ensure they have been met.

The development of the ML component for the ACTIONS commercial application was also described, and the key differences in applying the AMLAS process to the space and ground segments discussed.

3.4.3. Using Simulation (2.7)

Practical guidance was contributed for implementing simulation for the small space system domain. The guidance describes how simulation was implemented in ACTIONS for demonstration and HIL testing of the emergency response service application, and the creation of data products for ground segment applications.

3.4.4. Defining Understanding Requirements (2.2.1.2)

Practical guidance was contributed for defining the requirements for the 'Understand' element of the SUDA architecture of an autonomous system. The ML component of the ACTIONS emergency response service demonstration scenario was used as an example to describe defining requirements for this element.

4. IMPACT

4.1. BoK Alignment

Across the project consortium, work carried out by Craft Prospect and GSI has benefited from alignment to the AAIP Body of Knowledge. Alignment with AMLAS, requirements definition, SUDA architecture definition and simulation processes outlined in the BoK framework have had a positive impact on a range of current and future projects.

For Craft Prospect, positive impacts include:

- An established approach for assuring machine learning components of onboard operations across future projects.
- An established reputation for leading in assured ML, with published content reaching a wide audience.
- A developed simulation test bench and assured processing chain.
- Deployment of customer and partner hardware in simulation testing process.
- Establishment of an ongoing relationship with University of York, the Assuring Autonomy International Programme and mission end users.

For GSI, positive impacts include process improvements for ML-based products and services, such as:

- Standardisation of procedures and documentation while allowing innovation.
- Flow down of top-level requirements and increase of up-front work to benefit initial understanding of development data and suitable verification methods.

4.2. Engagements

Wider impacts include the presentation of work completed under the ACTIONS project to end users and industry professionals.

A workshop was held to present work to the CTO of the satellite platform and component manufacturer AAC Clyde Space. Positive feedback was provided on the project objectives and accomplishments, and potential future applications of the work discussed.

Presentation of work at past and upcoming conferences:

- **4S Symposium, Portugal – May 2022**
Oral presentation and paper
A bi-annual conference bringing together professionals from all over the world to engage in discussions about a wide variety of space topics. Technical sessions include mission and system analysis, Earth observation, science and new technologies.
- **ESA Living Planet Symposium, Germany – May 2022**
Oral presentation
An international Earth observation conference featuring many, wide ranging scientific sessions, where academics and industry professionals present their latest findings on

Earth's environment and climate derived from satellite data. The event focusses on the role of earth observation in building a sustainable future and a resilient society, and how business and the economy can benefit from emerging technologies.

- **UK National Earth Observation Conference, UK – September 2022**

Accepted for oral presentation and paper

A national Earth observation conference attracting the UK's EO and photogrammetry community across research, government and industry. The conference topic is 'Earth Observation Science - Technology in Action' and technical sessions are designed to cover all aspects of Earth observation.

ASSURING
AUTONOMY
INTERNATIONAL PROGRAMME