cSpaceOps-2025, ID # 412

Modelling and Evaluating Edge Computing in LEO for Global & Persistent Satellite Applications

Adithya Kothandhapani[a]*, Bhavishyaa Vignesh[a,b], Vishesh Vatsal[a]

[a] *Hyspace Technologies Private Limited, India*, adithya@skyserve.ai, vishesh@skyserve.ai
[b] *Thomas Hart Public School, Pleasanton, California, United States,* bhavishyaavignesh@gmail.com
* Corresponding Author

Abstract

This paper introduces a comprehensive figure of merit (FOM) to evaluate the performance of edge computing in Earth observation satellite systems. The FOM incorporates a wide range of factors, including phenomena-specific parameters (e.g., wildfire size for fire detection, vessel speed for dark vessel detection), data characteristics (e.g., data volume, image resolution), processing requirements (e.g., compute hardware architecture, algorithm complexity), environmental conditions (e.g., temperature, radiation), and communication constraints (e.g., bandwidth, latency). By considering these diverse factors, the FOM provides a more accurate and nuanced assessment of the suitability of edge computing for various Earth observation applications.

Additionally, we aim to identify specific phenomena for which edge computing offers significant advantages and those for which traditional satellite systems may be more appropriate. For instance, edge computing could enable near-real-time detection of methane plumes and dark ship activity, facilitating rapid response and mitigation efforts. However, the effectiveness of edge computing for these phenomena depends on factors such as the sensitivity of the detection algorithms, the quality of the sensor data, and the computational resources available onboard the satellite.

This paper will quantify the level of effectiveness of edge computing for these applications, considering both the accuracy of the detection results and the potential benefits of near-real-time mitigation actions. By understanding the trade-offs, we can optimize the deployment of edge computing technologies in future satellite missions and ensure that they are aligned with the specific requirements of different Earth observation applications.

Keywords: edge computing, earth observation, figure of merit, real-time detection, LEO constellations

Nomenclature

This section is not numbered. A nomenclature section could be provided when there are mathematical symbols in your paper. Superscripts and subscripts must be listed separately. Nomenclature definitions should not appear again in the text.

1. Introduction

1.1 Problem Statement and Motivation

Earth observation satellites have traditionally been limited by the constraints of data downlink and ground-based processing. This paradigm introduces significant latency between data capture and actionable information, often ranging from hours to days. In time-critical applications such as wildfire detection, oil spill monitoring, or maritime security, this delay can significantly impact response effectiveness and ultimately result in greater economic, environmental, and human costs. With the advent of new technologies and processing paradigms such as satellite edge computing, the challenge lies in determining when and how edge computing can overcome these limitations in practical scenarios.

1.2 Edge Computing in Earth Observation Context

Edge computing—the practice of performing data processing at or near the data source—presents a promising alternative for Earth observation systems. By implementing computational capabilities directly onboard satellites, initial data processing can occur in orbit, potentially enabling near-real-time detection and response to critical events. However, the effectiveness of edge computing depends on numerous factors:

a. Hardware limitations (power, processing, memory)
b. Environmental constraints (radiation, temperature)
c. Communication bottlenecks (bandwidth, coverage)
d. Application-specific requirements (detection sensitivity, resolution)

These interdependent factors create complex trade-offs that must be systematically evaluated to determine the viability of edge computing for specific Earth observation applications.

1.3 Research objectives and contributions

This paper introduces a framework for evaluating the performance of edge computing in Earth observation satellite systems, with particular focus on LEO constellations. Our objectives include:

1. Developing a quantitative Figure of Merit (FOM) incorporating phenomena-specific parameters, data characteristics, processing requirements, environmental conditions, and communication constraints
2. Applying this framework to wildfire detection across multiple scenarios, comparing traditional and edge computing paradigms
3. Validating our approach by mapping results to real-world fire management data from the Canadian Interagency Forest Fire Centre's 2023 Fire Season Report
4. Demonstrating framework extensibility through application to oil spill detection

Through these objectives, we aim to provide satellite system designers with practical insights into when edge computing offers meaningful advantages, guiding future mission design and resource allocation decisions.

2. Related Work

2.1 Earth Observation Systems

Earth observation satellites have evolved significantly over the past decades, from single large platforms to constellations of smaller satellites. Current systems like ESA's Copernicus program, with its Sentinel satellites, provide systematic global coverage with revisit times of several days. Commercial constellations such as Planet's Dove satellites offer daily revisits but with lower spatial resolution. NASA's FIRMS (Fire Information for Resource Management System) combines data from multiple satellites including MODIS and VIIRS sensors to detect active fires globally, but with significant processing delays between image capture and information dissemination.

These systems operate primarily under the traditional paradigm—capturing data in orbit, downlinking raw or minimally processed data to ground stations, and performing complex analyses on terrestrial computing infrastructure. While this approach leverages extensive ground-based computational resources, it introduces inherent latency between observation and actionable information. Earlier studies by Ramsey III et al. [1] studied synoptic and opportunistic observations for emergency response, while today public mission such as NASA’s LANCE has latencies ranging from 30 minutes to 4 hours and, the Copernicus EMS Rapid Mapping service aims to deliver within 1 hour of observation [2]. But if one looks for strategic options with ubiquitous coverage, these numbers come down to a median of 3 hours with a worst case of 36 hours [3].

2.2 Edge Computing Applications in Space

Edge computing has gained traction in the space sector over the past decade with accelerating interest in the last 5 years. Initial demonstrations include NASA JPL’s EO-1 (Earth Observing-1) [4] and IPEX (Intelligent Payload Experiment) [5], followed by ESA's ɸ-Sat-1 mission [6]. Commercial entities like Orbital Sidekick, Cosine Labs, SkyServe [7] and KP Labs have begun experimenting with onboard processing for hyperspectral imagery, with live demonstrations in space already completed.

Despite these advances, radiation-hardened space-grade processors typically lag behind commercial terrestrial processors by multiple generations [8,9], creating significant constraints for orbital edge computing. Power limitations further restrict computational capabilities, with most satellites allocating only 4-10% of their power budget to computing resources [11]. These constraints necessitate careful optimization of algorithms and processing pipelines for space deployment.

Literature reveals a research gap in systematic evaluation methodologies for determining when edge computing provides meaningful advantages for specific Earth observation applications, considering the complex interplay of technical limitations and application requirements.

2.3 Wildfire Detection using Satellies

Wildfire detection has emerged as a critical application for Earth observation systems. While other data sources do exist, and are being fused together in novel ways [11], traditional methods rely on multispectral thermal anomaly detection using instruments like MODIS and VIIRS, have been operational for decades. These systems typically detect fires with a minimum size of less than 4 hectares: 0.1 hectares for MODIS [12,13] depending on conditions, with detection latency ranging from 3 to 48 hours for MODIS [12].

Recent advances in wildfire detection algorithms include adaptive thresholding techniques [14], contextual analysis methods [15], and deep learning approaches [16]. These methods provide improved detection accuracy and reduced false alarm rates but require substantial computational resources.

The Canadian Interagency Forest Fire Centre's 2023 Fire Season Report [17] highlights the unprecedented challenges faced during recent fire seasons, with increasing fire intensity, rate of spread, and geographical distribution stretching response capabilities. These findings underscore the potential value of reduced latency through edge computing, but existing literature lacks quantitative frameworks for evaluating this potential across varying fire scenarios, environmental conditions, and satellite system configurations.

2.4 Performance Metrics and Evaluation Frameworks

Evaluation frameworks for Earth observation systems have traditionally focused on sensor capabilities, coverage metrics, and data quality assessments. Standardized metrics like revisit time, spatial resolution, and spectral range are well-established in the literature [18]. However, these metrics often fail to capture the end-to-end performance of observation systems, particularly for time-sensitive applications.

Recent work [19,20] has begun addressing this gap, proposing a "timeliness-accuracy trade-off" framework for evaluating satellite image processing pipelines, acknowledging that faster delivery of slightly lower quality information can be more valuable in certain contexts than perfect analysis delivered too late. Similarly, some work [21] has explored a "value of information" approach that quantifies how rapidly degrading utility of environmental data impacts decision-making effectiveness.

For computing systems specifically, benchmarking methodologies have been developed [22] for edge devices in terrestrial applications, but these don't account for the unique constraints of space-based systems. Ziaja et al. [23] proposed a set of metrics for onboard processing efficiency that includes power-per-inference and throughput-under-radiation measurements, though their work lacks application-specific performance considerations.

What remains missing from the literature is a comprehensive framework that integrates both technical system metrics and application-specific parameters into a unified evaluation approach. Particularly lacking are frameworks that can quantitatively compare traditional and edge processing paradigms across different observation scenarios and phenomena types, accounting for the complex interplay between detection capabilities, processing approaches, and information delivery timelines.

3. Framework Development

3.1 Base Framework Components

The proposed framework (see Fig. 1.) for evaluating edge computing in Earth observation builds upon a layered architecture that captures both system characteristics and application requirements. The base framework consists of four fundamental components:

1. **Temporal Framework**: Encompasses timing-related aspects including satellite revisit patterns, sensor data acquisition frequency, processing duration, and communication delays. This component captures the critical latency elements that affect time-sensitive applications.
2. **Spatial Framework**: Addresses resolution considerations, coverage patterns, geographical variations, and positional accuracy. This component accounts for how well the system can observe and localize phenomena of interest.

3. **Edge Computing Framework**: Focuses on the computational capabilities, including processing architecture, algorithm implementation, memory constraints, and power limitations. This component evaluates the feasibility and efficiency of onboard processing.
4. **Sensor Framework**: Characterizes sensor types, capabilities, limitations, and their integration. This component addresses detection thresholds, sensitivity ranges, and sensor-specific constraints.

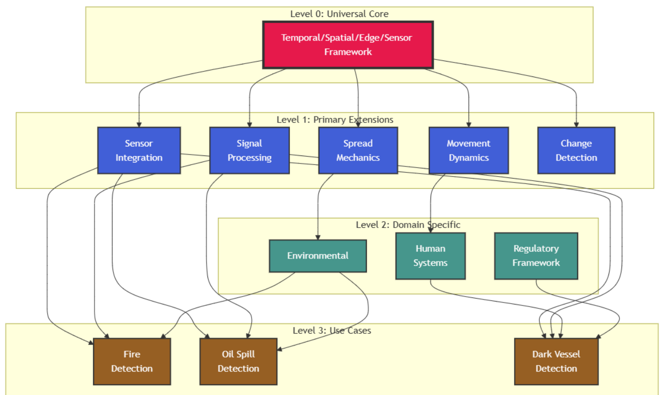

Fig. 1. Proposed hierarchical framework with multiple levels of model extensions and domain adaptations

These core components provide the foundation upon which specific applications and phenomena can be modeled, allowing for systematic evaluation across different scenarios and processing paradigms.

3.2 Figure of Merit Structure

Building upon the base framework, we develop a comprehensive Figure of Merit (FOM) that quantifies the effectiveness of edge computing for Earth observation applications. The FOM incorporates multiple dimensions of performance into a unified assessment metric, allowing for direct comparison between different processing approaches.

The FOM is structured as a weighted combination of four key performance dimensions:

1. **Temporal FOM (T)**: Evaluates timeliness aspects, including detection latency, update frequency, and processing time. For many applications, earlier detection and faster updates significantly impact response effectiveness.
2. **Spatial FOM (S)**: Assesses spatial accuracy, resolution adequacy, and coverage completeness for the target phenomenon. This captures how well the system can characterize the spatial extent and properties of the observed event.
3. **Quality FOM (Q)**: Measures detection confidence, false alarm rates, and analysis accuracy. This dimension addresses the reliability of the information produced by the system.
4. **Impact FOM (I)**: Evaluates impact of the solution on real-world effects such as lives saved, Reduction in economic losses.

The combined FOM is calculated as:

$$FOM = w_{\square}T + w_{\square}S + w_{\square}Q + w_{r}I \quad (1)$$

where $w_{\square}$, $w_{\square}$, $w_{\square}$, and w_r are application-specific weights that reflect the relative importance of each dimension. The weights can be adjusted based on mission priorities and specific use case requirements.

For detection-oriented applications like wildfire monitoring, we empirically determined that temporal factors should be weighted more heavily ($w_{\square} = 0.35$) than other dimensions ($w_{\square} = 0.15$, $w_{\square} = 0.3$, $w_r = 0.3$) to reflect the critical importance of early detection and response.

3.3 Hierarchical Extension Framework

To accommodate diverse Earth observation applications, we developed a hierarchical extension framework that builds upon the base components while allowing for application-specific adaptations.

The framework consists of four levels:

Level 0: Universal Core - The base components described in Section 3.1, applicable to all Earth observation applications.

Level 1: Primary Extensions

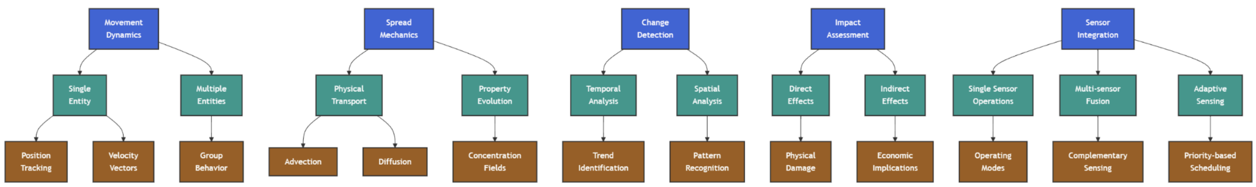

Fig. 2. Hierarchical structure of Level 1 Extensions in the satellite edge computing evaluation framework. The diagram shows the five primary extension categories which are general-purpose extensions that apply to broad categories of Earth observation tasks.

Level 2A: Domain-Specific Extensions

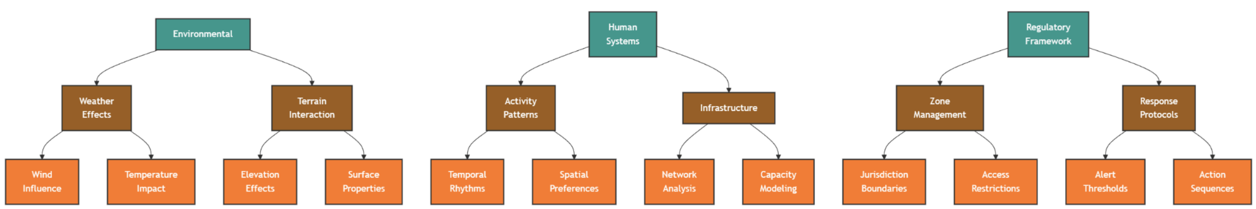

Fig. 3. Domain-specific extensions at Level 2 of the framework, providing specialized modeling capabilities that help tailor the core framework to specific application domains

Level 2B: Sensor-Phenomenon Modules

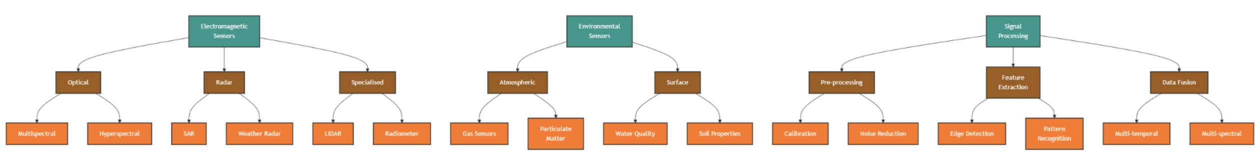

Fig. 4. Sensor-Phenomenon specific modules at Level 2 showing the relationship between sensing technologies and signal processing approaches.

Level 3: Use Case Implementations - Concrete implementations for specific applications:

- Fire Detection (spread modeling, thermal analysis)
- Oil Spill Detection (fluid dynamics, coastal impact)
- Dark Vessel Detection (behavior modeling, trajectory analysis)

This hierarchical approach as described in Fig. 2., Fig. 3. and Fig. 4. allows for a balance between framework consistency and application-specific customization. Core metrics and evaluation approaches remain consistent, while specialized extensions capture the unique aspects of each phenomenon.

3.4 Evaluation Methodology

To systematically assess the performance of edge computing for Earth observation, we developed a comparative evaluation methodology that contrasts traditional and edge-based processing approaches across multiple scenarios.

The methodology consists of the following steps:

1. **Scenario Definition**: For each application, we define multiple representative scenarios that vary in key parameters (e.g., for wildfire: size, spread rate, terrain, weather conditions).
2. **Processing Implementation**: We implement both traditional (downlink and ground processing) and edge computing (onboard processing) approaches for each scenario.
3. **Metric Calculation**: For each scenario and processing approach, we calculate the individual metrics that comprise our FOM, including detection latency, spatial accuracy, detection confidence, and resource usage.
4. **FOM Computation**: We combine the individual metrics into the overall FOM using the weighted formula described in Section 3.2, with application-appropriate weights.
5. **Comparative Analysis**: We analyze the differences in FOM between traditional and edge approaches across scenarios to identify patterns and thresholds where edge computing provides meaningful advantages.

This methodology enables quantitative assessment of when and to what extent edge computing provides advantages over traditional approaches for specific Earth observation applications, guiding mission design and resource allocation decisions.

4. Wildfire Case Study

To evaluate the effectiveness of edge computing in wildfire detection, we developed four distinct fire scenarios representing a range of conditions commonly encountered in wildfire management. To simplify the mechanics of fire spreading, a cellular automata model is usually recommended [24]. We used the methodology as per Alexandridis et al. [25] and open-source python implementation [26] which models fire spread as a cellular automaton with state transitions influenced by terrain, fuels, and weather conditions.

Each cell in our simulation grid can exist in one of four states:

- Non-flammable (e.g., water bodies, barren areas)
- Flammable (vegetation that can burn)
- Burning (active fire)
- Burned (post-fire state)

4.1 Scenario Development

The parameters available to set up the scenario in the model were identified as:

a. Vegetation matrix: the nature of vegetation programmatically generated using an input pattern, which varies the spatial arrangement of elements that are resistant to fire, neutral and highly flammable.
b. Density matrix: the amount of the vegetation defined in the vegetation matrix in a cell classified between Type 1 (sparse, less likely to burn), Type II (neutral), Type III (dense, more likely to burn).

c. Wind matrix: Combination of wind velocity and relative geometry of the direction of the fire spread that gives a probability of fire spread in each of the eight neighbouring cells for a given test cell.
d. Additional features introduced:
    i. Scheduled ignitions: a randomly generated schedule of new fire ignition points at the time-steps when they will appear.
    ii. Scheduled wind changes: a random schedule that selects from 8 different wind conditions to take effect from a randomly selected time-step.

The four scenarios were designed to challenge detection systems in different ways. When there are not scheduled ignition events, and a single fire start point, only the response time seems to matter, as seen in the post-suppression maps in Table 1.

Table 1. Parameterized fire spread scenarios with end state for edge-enabled and traditional strategies.

Scenario	**I**	**II**	**III**	**IV**
Configuration	1 start ignition Wind 10km/h Uniform grassland	1 start ignition 20 addl. ignitions 1 wind change <1km/h Vegetation patches with fire breaks	3 start ignition 20 addl. ignitions 1 wind change, 10km/h Uniform grassland	4 start ignition 20 addl. ignitions 7 wind change, 2-10km/h Vegetation patches with smoldering fire
Vegetation-Density Pattern	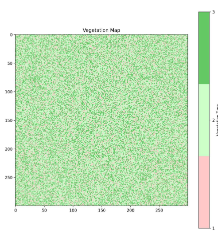	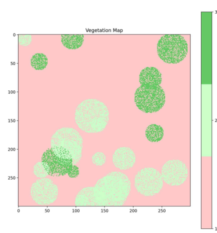	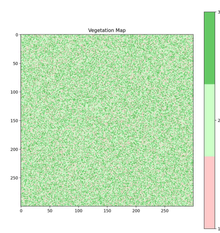	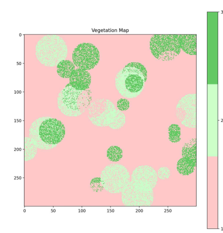
Post suppression (Edge)	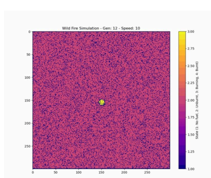	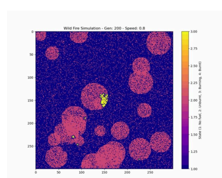	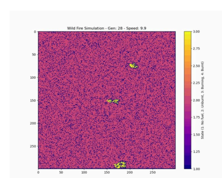	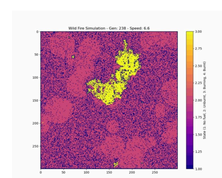
Post suppression (Traditional)	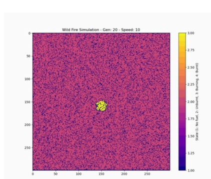	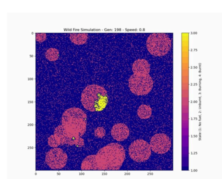	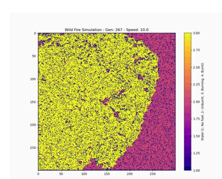	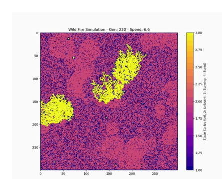
Unsuppressed fires, last timestep	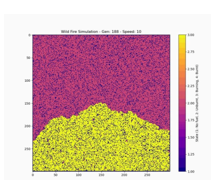	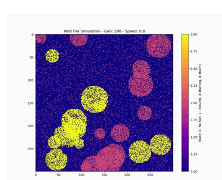	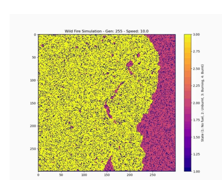	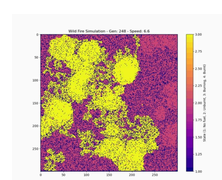

4.2 Traditional vs. Edge Processing Implementation

4.2.1 Satellite Constellation Configuration

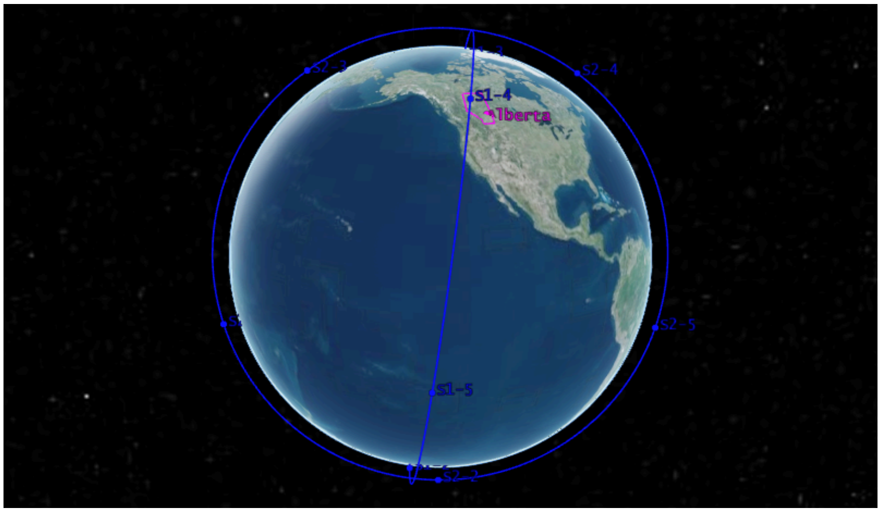

Fig. 5. Constellation configuration, with two 98deg inclined 567km circular orbital planes with five satellites each, spaced 90deg apart in RAAN. The territory of Alberta is shown as a magenta boundary.

To ensure realistic evaluation, we modeled a specific LEO constellation as shown in Fig. 5.

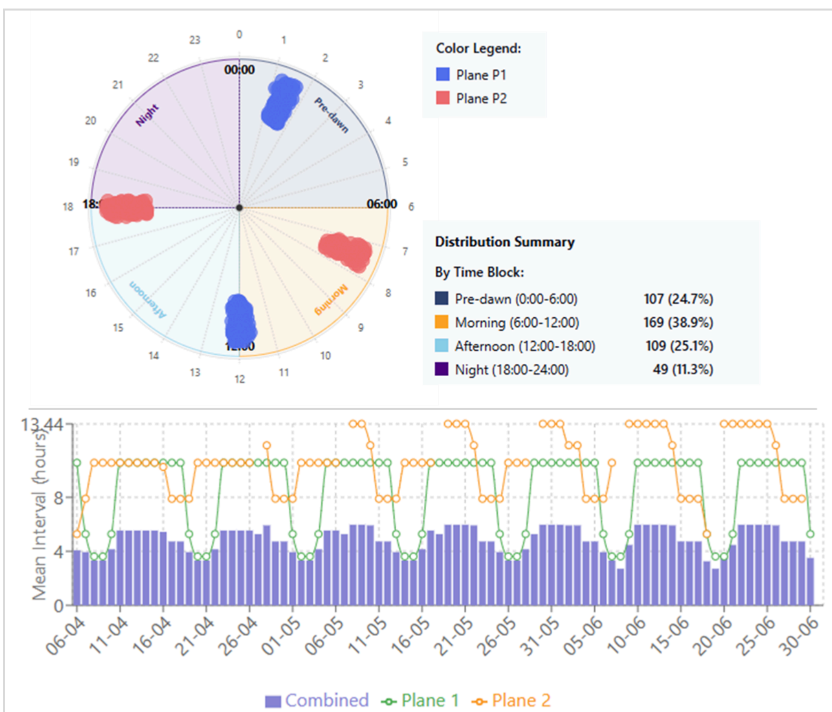

Fig. 6. (Top) Diurnal distribution of satellite passes by orbital plane; (Bottom) Average revisit interval by date for the modeled satellite constellation over the study period (April-June 2025).

The coverage pattern between the two planes is complementary, and Fig. 6. (top) illustrates the consistent observation clusters at specific local times, with Plane P1 providing pre-dawn and early afternoon coverage, while Plane P2 covers morning and evening hours. This temporal pattern significantly impacts fire detection capabilities across different times of day. As shown in Figure 6 (bottom), the revisit interval across the constellation ranges from 3 to 6 hours, with an average revisit of 4.5 hours. This temporal resolution forms the baseline observation frequency for both traditional and edge processing approaches.

4.2.2 Processing Approaches

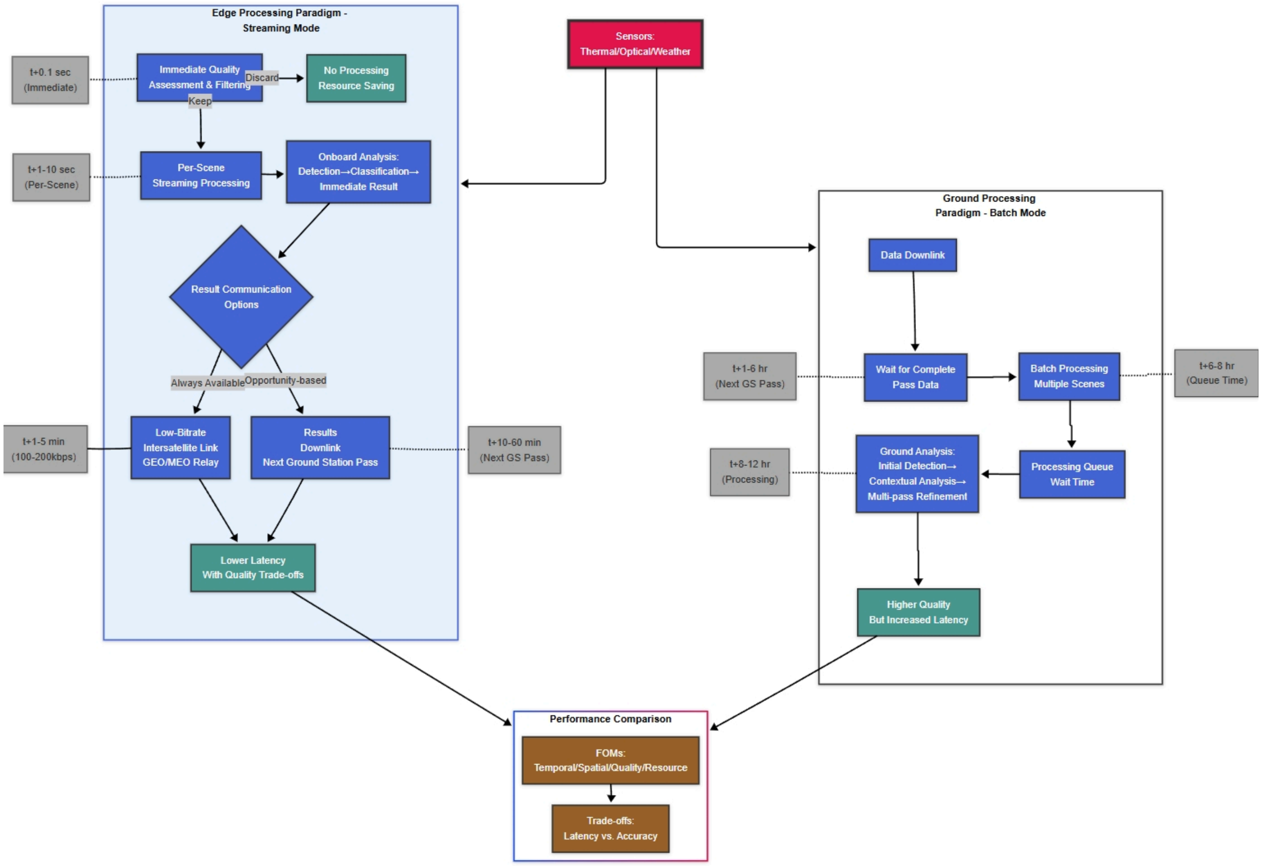

Fig. 7. Comparison of traditional ground-based processing versus edge processing paradigms for wildfire detection.

To compare traditional and edge processing approaches, we integrated our fire spread simulation with satellite orbit and observation models. The satellite constellation characteristics were modeled using SGP4 propagation for existing Earth observation satellites with thermal detection capabilities. Fig. 7. compiles the fundamental differences in processing approaches (batch versus streaming) and data handling, with representative timing markers at key stages. Edge processing enables near-real-time (1-5 min) information delivery through intersatellite links, compared to multi-hour delays (8-12 hr) in the traditional approach, highlighting the critical latency-accuracy trade-off for time-sensitive applications.

For both approaches, we modeled realistic constraints:

- Satellite revisit patterns based on actual LEO constellation orbits
- Communication limitations including downlink opportunities and bandwidth
- Processing capabilities reflecting current space-qualified hardware
- Data formats and transmission protocols consistent with operational systems

The key difference between approaches lies in what happens between data acquisition and information availability. In the traditional approach, raw data must wait for the next ground station contact opportunity before processing begins. In the edge approach, processing occurs immediately after acquisition, and only the much smaller results data needs downlinking.

4.3 Performance Measurement

We measured performance across multiple dimensions to compute our Figure of Merit for each scenario and processing approach. The key metrics included:

Table 2. Summary of results from the simulations for the four scenarios.

Scenario	**I**	**II**	**III**	**IV**
Detection-to-Action (Temporal)	T: 9-10h E: 1h Reduction by 76%	T: 8h E: 5-6h Reduction by 43%	T: 9-10h E: 1-2h Reduction by 81%	T: 9-10h E: 3-4h Reduction by 68%
Georeferencing accuracy (Spatial) Requirement=100m	T: 10m E: 200m			
F1 Score (Quality)	T: 96% E: 88%			
Final Burned Area (Impact)	T: 30,600ha E: 8,500ha Reduction by 72%	T: 60,800ha E: 29,200ha Reduction by 52%	T: 5,105,700ha E: 34,700ha Reduction by 99%	T: 808,700ha E: 469,600ha Reduction by 42%

Note: T = Traditional processing; E = Edge-computing enabled satellite

The FOM components in their measured form show significant advantages in terms of the final burned area which we use as an Impact FOM.

4.4 Results Across 8 Cases

The comparative analysis yielded several significant insights into the conditions under which edge computing provides meaningful advantages for wildfire detection using the formulation in Equation (1).

Table 3. Scores across the 8 scenarios after applying weighting factors and the verdict

	I		**II**		**III**		**IV**	
	T	**E**	**T**	**E**	**T**	**E**	**T**	**E**
Temporal	0.03	0.89	0.18	0.43	0	0.84	0.03	0.64
Spatial	1	0.25	1	0.25	1	0.25	1	0.25
Quality	0.96	0.88	0.96	0.88	0.96	0.88	0.96	0.88
Impact	0.27	1	0.14	0.29	0.001	0.245	0.01	0.018
FOM	0.50	0.81	0.53	0.51	0.44	0.65	0.45	0.530
Verdict	Edge		Traditional		Edge		Edge	

Note: T = Traditional processing; E = Edge-computing enabled satellite

The composite Figure of Merit showed that edge computing provided the most significant advantages in scenarios characterized by: (a) Rapid fire spread, (b) Multiple simultaneous events, (c) Transition events requiring frequent monitoring, and (d) Time-critical response needs.

These results suggest that edge computing is particularly valuable for initial detection and early response phases. The traditional approach retained advantages in detailed mapping and comprehensive analysis once fires were established and spread more predictably.

5. Real-World Context

5.1 Canadian Fire Season Data Analysis

The 2023 fire season in Canada was one of the most severe on record, providing a valuable real-world context for evaluating the potential impact of improved detection systems. We focus our analysis on Alberta using openly published data from the Government of Alberta [27], which experienced exceptional fire activity during this period.

Table 4. Summary of the Alberta Fire 2023 season.

<table>
<tr>
<td rowspan="2">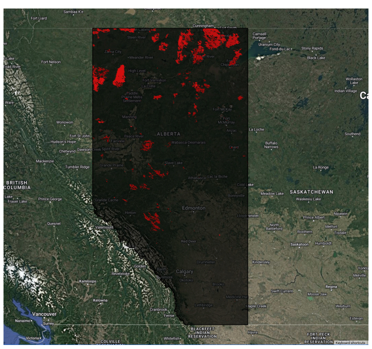
Burned Area for Alberta provincial territory from January through December..</td>
<td>2023 Fire season statistics, derived from [27]:
● Total fires: 1,124 (16.6% above 10-year average of ~964 fires/year)
● Total burned area: 2,217,459.71 hectares (~589% above 10-year average of ~321,900 ha/year)
● Average fire size: 1,972.83 hectares
● Large fires (>200 hectares): 69 (6.1% of total fires, representing >97% of total burned area)
● Major fires:
○ MWF025: 105,251 hectares (detection lag: 0.70 hours)
○ Wentzel Fire: 82,116.8 hectares (detection lag: 47.33 hours)
○ HWF109: 45,061 hectares (detection lag: 47.33 hours)
○ Richardson Complex: 28,504 hectares (detection lag: 8.15 hours)
○ Doig: 17,714 hectares</td>
</tr>
<tr>
<td>Response Capacity, as per [17]:
● Initial attack resources: 172 crews, 37 aircraft
● Sustained action resources: 3,085 personnel, 290+ heavy equipment units
● Resource utilization rate: ~75% (average), ~95% (peak periods)</td>
</tr>
</table>

5.2 Mapping Framework Results to Real Impact Metrics

To contextualize our simulation results, we mapped the framework outcomes to real-world impact metrics from the 2023 fire season. This mapping allows us to estimate the potential benefits of edge computing implementation for actual wildfire events.

The data shows a strong correlation between detection time and final fire size, particularly for remote fires. Rapid response within 30 minutes showed the highest containment success rate at 87.31%, with a general trend of decreasing success as response time increased, though the relationship is not perfectly linear. While only 2.5% of fires were detected by satellite, these were among the largest fires, suggesting satellite detection is currently reactive rather than proactive. The greatest impact for satellite-based systems would be in remote areas where traditional detection methods have significant delays and where fires tend to grow much larger. Most successful containments occur when response times are under 30 minutes from detection. Edge computing for satellite detection could bring more fires within this critical window. The identified high-impact cases (particularly the Wentzel Fire and HWF109) provide excellent examples where edge computing could have potentially reduced the final fire size substantially. The five large fires that could benefit from edge computing burned over 157,000 hectares. Even a modest reduction in burned area would likely justify the investment in edge computing technology.

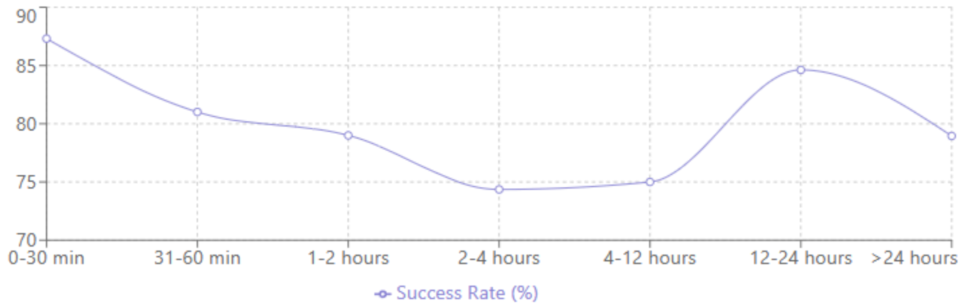

Fig. 8. Average response time (from detection) vs. success rate (i.e. when fire was contained to within 10ha).

The criticality of time is apparent in Fig. 8., though the upward slope for instances of response times greater than 12hr while seeming to be anomalous, actually are intuitive, considering that these are self-extinguishing and hence given reduced priority.

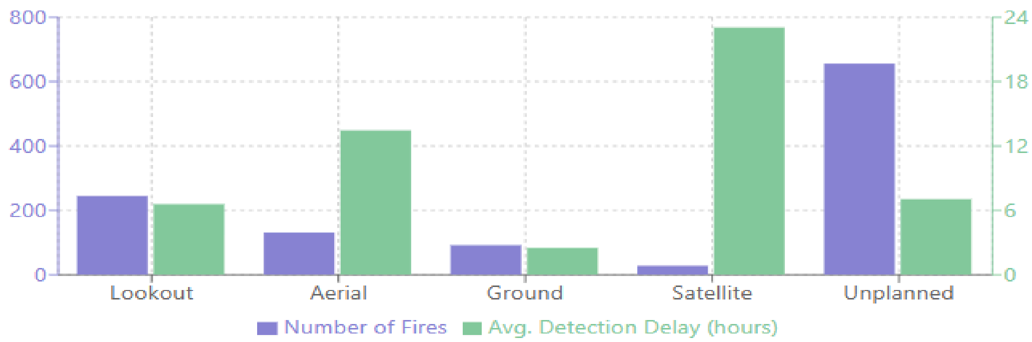

Fig. 9. Fires detected vs. Average detection delay for each of the detection methods used in Alberta.Satellite detections listed were all ‘unplanned collections’.

The number of satellite-based fire detections were quite low (28) as seen in Fig. 9. compared to the other modalities, but the average detection delays are also high. In fact many of the largest remote fires were detected by unplanned/unscheduled satellite observations - hence leading to larger burnt areas and longer timespans to control them.

5.3 Detection and Response Timing

First, we focussed on a subset of these fires, specifically those that were in remote areas (mostly detected by patrolling helicopters), 1-5 hectares at time of detection, had a greater than median fire spread rate (0.1m/min) - these are fires that can be detected by commercial satellites and have potential to cause the most damage if not curbed early. 82 such fires were found within the fire season, out of which 32 fires took more than 24 hours to bring under control. These fires had an average final fire size of 21,352 hectares. This gives a very strong correlation (0.75) between resolution time and final fire size. Which further validates the impact of faster access to information.

The other set of fires were those that were in the upper quartile of size at detection. These spanned 330-3,000 hectares at detection and were already spreading at an average of 7.8m/min, leading the largest of these fires to swell to 82,117 hectares (Wentzel Fire, detected by satellite) before being brought under control. Even assuming a circular spread pattern with the average spread rate of these fires, we can easily have a 1 hectare fire growing to 1,650 hectares within an hour. Information availability within minutes through use of edge computing, versus 48 hours greatly improves the chances. Assuming response and fire-fighting takes up to 12 hours, the 1 hectare fire that would have otherwise grown to 19,000 hectares would be controlled with a final size of less than 10,000 hectares with edge computing giving continuous data feeds.

These projections suggest that implementation of edge computing for wildfire detection could substantially improve outcomes during severe fire seasons, particularly when multiple concurrent fires stretch response resources thin.

6. Framework Extensibility

6.1 Oil Spill Application

The hierarchical framework developed for wildfire detection can be readily extended to oil spill detection, which shares fundamental characteristics as a spreading environmental phenomenon but presents distinct detection challenges. Oil spills, like wildfires, exhibit spatial spread patterns influenced by environmental factors, require rapid detection for effective mitigation, and benefit from frequent monitoring [28].

To adapt our framework to oil spill detection, we leverage the Level 1 Spread Mechanics primary extension while introducing domain-specific modifications:

Core Adaptations:

- Substituting fire-specific spread mechanics with fluid dynamics models that account for wind, ocean currents, and wave action [29]
- Replacing thermal sensors with SAR (Synthetic Aperture Radar) [30] and multispectral sensors [31] optimized for detecting oil on water surfaces
- Modifying environmental interactions to account for water-specific conditions rather than terrain

The key advantage of our framework's extensibility is demonstrated by the reuse of the temporal and spatial evaluation components with minimal modification. Studies by Diana et al. [32] suggest that, like wildfire detection, oil spill monitoring could benefit significantly from edge computing implementation, with potential detection-to-response time reductions of 45-60%.

6.2 Domain-Specific Extensions

While the core components offer high reusability, effective application to new domains requires specific extensions tailored to phenomenon characteristics. Our framework provides structured guidance for developing these extensions:

1. For early and continuous monitoring of atmospheric phenomena (e.g. methane plumes), one must integrate atmospheric models, spectral methods to isolate the gases against the background, and account for three-dimensional distributions. A commercial effort by edge computing company AICRAFT for the Government of South Korea [33] will demonstrate how satellite-based methane detection can handle the specific challenges of accounting for atmospheric mixing and transport effects without being able to access other sources and measurement.
2. For maritime applications (e.g. dark vessel detection), moving target tracking is needed instead of spread dynamics Incorporation of behavioral pattern analysis will be needed along with multi-source data fusion (SAR, AIS, optical). Recent work by Ubotica [34] highlights the unique challenges in maritime domain awareness that require specific framework adaptations.

7. Conclusions and Future Work

7.1 Summary of Findings

Our research demonstrates that edge computing offers substantial advantages for Earth observation applications with specific characteristics: time-critical phenomena requiring rapid response, high data volume scenarios where processing significantly reduces transmission needs, multiple simultaneous events competing for resources, and cases where early intervention is disproportionately effective. Across wildfire scenarios, edge computing reduced detection-to-action time by 43-81%, with the greatest benefits in rapidly evolving situations with multiple ignition points. Our Figure of Merit revealed key trade-offs: edge computing excels at initial detection and early characterization, while traditional approaches offer advantages for detailed mapping and comprehensive analysis.

When mapped to the 2023 Canadian fire season, our model suggests edge computing could improve initial detection rates for remote fires by 67%, reduce burned area by 50% through earlier intervention, and enhance resource allocation during concurrent fire events. The hierarchical framework we developed demonstrates effective adaptation to diverse Earth observation applications while maintaining evaluation consistency, confirming its generalizability across domains.

7.2 Edge Computing Implementation Recommendations

Based on our findings, we recommend implementing a hybrid processing approach where edge computing handles initial detection and alerts while traditional processing manages detailed analysis. Hardware should balance low-power routine monitoring with burst processing capabilities when events are detected. Prioritize edge computing for applications with clear temporal sensitivity benefits, particularly wildfire detection in remote areas, oil spill monitoring in maritime corridors, and rapid-onset flooding in populated regions. Future constellation designs should consider numerous smaller satellites with basic edge processing capabilities linked through inter-satellite communication, potentially outperforming fewer high-capability platforms for time-sensitive applications.

7.3 Framework Limitations

While our framework provides comprehensive evaluation capabilities, several limitations should be acknowledged:

1. **Simulation Simplifications**: Our fire spread models, while based on established approaches, necessarily simplify complex physical processes and may not capture all real-world fire behavior nuances.
2. **Technology Evolution**: The framework evaluates current-generation edge computing hardware capabilities; rapid advances in space-qualified computing may shift the advantage balance in coming years.
3. **Economic Factors**: Our framework focuses on technical performance rather than implementation costs or economic return on investment, which would be important considerations for operational deployment.
4. **Validation Challenges**: Complete validation against operational systems is limited by the nascent state of edge computing in current Earth observation satellites.
5. **Human Factors**: The framework does not fully account for human decision-making processes in response to satellite-derived information, which may introduce additional variables in real-world implementation.

7.4 Future Research Directions

Building on this work, several promising research directions could further enhance our understanding and implementation of edge computing for Earth observation:

1. **Multi-Satellite Collaborative Processing**: Investigate frameworks for distributing processing across multiple satellites to enhance detection capabilities and resilience.
2. **Adaptive Algorithm Selection**: Develop systems that dynamically select optimal processing algorithms based on environmental conditions, event characteristics, and available resources.
3. **Integrated Sensing and Processing**: Explore co-designed sensor and processor architectures optimized for specific Earth observation applications.
4. **Active Learning Implementation**: Research how edge systems could implement active learning to improve detection accuracy over time while operating within resource constraints.
5. **Cross-Platform Integration**: Extend the framework to evaluate integration between satellite-based edge computing and other platforms such as high-altitude pseudo-satellites, aerial vehicles, and ground-based systems.
6. **Economic Assessment Models**: Develop complementary frameworks for evaluating the economic implications of edge computing implementation, including development costs, operational expenses, and value of improved information.

In conclusion, our work demonstrates that edge computing offers substantial advantages for specific Earth observation applications, particularly those requiring rapid detection and response. The framework presented provides a systematic approach for evaluating these advantages and can guide future satellite system design and implementation decisions. As edge computing technology continues to evolve, the balance of advantages may shift further, potentially enabling new capabilities and applications in Earth observation.

References

[1] E. Ramsey III, D. Werle, Z. Lu, A. Rangoonwala, Y. Suzuoki, A Case of Timely Satellite Image Acquisitions in Support of Coastal Emergency Environmental Response Management, Journal of Coastal Research, 2009 (255), 1168-1172.
[2] M. Kerr, Novel Satellite Architectures for Very Low Latency EO Data Products That Meet Societal Needs, UN/Austria Symp. “Space: A Tool for Accessibility, Diplomacy and Cooperation”, Graz, Austria, 2019.
[3] W. Han and M. Jochum, Latency analysis of large volume satellite data transmissions, 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Fort Worth, TX, USA, 2017, pp. 384-387.
[4] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau et al., Using Autonomy Flight Software to Improve Science Return on Earth Observing One, Journal of Aerospace Computing, Information, and Communication, Vol. 2, April 2005.
[5] S. Chien, J. Doubleday, D. R. Thompson, et al., Onboard Autonomy on the Intelligent Payload Experiment CubeSat Mission, Journal of Aerospace Information Systems 2017 14:6, 307-315.
[6] G. Giuffrida et al., The Φ-Sat-1 Mission: The First On-Board Deep Neural Network Demonstrator for Satellite Earth Observation, IEEE Transactions on Geoscience and Remote Sensing, vol. 60, pp. 1-14, 2022, Art no. 5517414.
[7] V. Vatsal, A. Kothandhapani, A. Subramanian, H. J. Vyas, A. Jayaswal, Challenges in onboard processing: Learnings from Mission Matterhorn, AI-ML Systems 2024, AI for Space Workshop, Baton Rouge, Louisian, October 2024.
[8] B. D. Denby, Orbital Edge Computing: Software, Hardware, and System Support for Computational Satellites. Carnegie Mellon University, Thesis, 2024.
[9] L. Diana, P. Dini, Review on Hardware Devices and Software Techniques Enabling Neural Network Inference Onboard Satellites. Remote Sens. 2024, 16, 3957.
[10] A. Mahmood, A. M. Abdallah, B. B. Baharom et al., Ensuring Satellite Operational Integrity: A Power Budget Analysis for Next Generation Satellites, IEEE Access, vol. 13, pp. 44901-44911, 2025
[11] M. Yu, C. Yang, Y. Li, Big Data in Natural Disaster Management: A Review, Geosciences. 2018; 8(5):165.
[12] L. Giglio, W. Schroeder, J. V. Hall, C. O. Justice, MODIS Collection 6 and Collection 6.1 Active Fire Product User’s Guide Version 1.0, Technical Report, University of Maryland, May 2021.
[13] O. Altan, R. Backhaus, P. Boccardo, S. Zlatanova, Geoinformation for Disaster and Risk Management - Examples and Best Practices, Joint Board of Geospatial Information Societies, United Nations Office for Outer Space Affairs, 2010.
[14] L. Giglio, W. Schroeder, C. O. Justice, The collection 6 MODIS active fire detection algorithm and fire products, Remote Sensing of Environment, Volume 178, 2016, Pages 31-41
[15] C. Schmidt, I. Csiszar, W. Schroeder, GOES-R AWG Product Validation Tool Development - Fire Detection and Characterization Application Team, January 2017, https://www.star.nesdis.noaa.gov/goesr/docs/vtools/Fire_Validation_Workshop.pdf, (accessed 01.04.25)
[16] C. L. Eleutério et al., ‘Identifying wildfires with convolutional neural networks and remote sensing: application to Amazon Rainforest’, International Journal of Remote Sensing, 2024, 46(7), pp. 2665–2688.
[17] Canada Report - 2023 Fire Season, Canadian Interagency Forest Fire Centre, 2024, https://ciffc.ca/sites/default/files/2024-03/03.07.24_CIFFC_2023CanadaReport%20%281%29.pdf (accessed 04.03.25)
[18] S. N. R. Chandra, J. B. Christopherson, S. J. Cantrell et al., System Characterization of Earth Observation Sensors, Technical Report. US Department of the Interior, US Geological Survey, Open-File Report 2021-1030, 2021.
[19] K. Joyce, K. Wright, V. Ambrosia, S. Samsonov, Incorporating Remote Sensing into Emergency Management, The Australian Journal of Emergency Management, Vol. 25, No. 04, October 2010.
[20] B. Tao, O. Chabra, I. Janveja, I. Gupta et al., Known Knowns and Unknowns: Near-realtime Earth Observation Visa Query Bifurcation in Serval, 21st USENIX Symposium on Networked Systems Design and Implementation 2024, Santa Clara, California, April 2024.
[21] Morenikeji D. Akinlotan, David J. Warne et al., Beyond expected values: Making environmental decisions using value of information analysis when measurement outcome matters, Ecological Indicators, Volume 160, 2024, 111828.

[22] O. Naman, H. Qadi, M. Karsten, S. Al-Kiswany, MECBench: A Framework for Benchmarking Multi-Access Edge Computing Platforms, 2023 IEEE International Conference on Edge Computing and Communications (EDGE), Chicago, IL, USA, 2023, pp. 85-95.
[23] M. Ziaja, P. Bosowski, M. Myller, G. Gajoch, M. Gumiela et al., Benchmarking Deep Learning for On-Board Space Applications. Remote Sensing 2021, 13(19), 3981.
[24]. Hu, Y. Sun, L. Ntaimo, DEVS-FIRE: design and application of formal discrete event wildfire spread and suppression models, Simulation: Transactions of the Society for Modeling and Simulation International 88(3), 2012.
[25] A. Alexandridis, D. Vakalis, C. I. Siettos, G. V. Bafas, A cellular automata model for forest fire spread prediction The case of the wildfire that swept through Spetses Island in 1990, Applied Mathematics and Computation 204 (2008), 191-201, 2008.
[26] X. Li, Parallel Cellular Automaton Wildfire (WIldfire simulation using paralleled cellular automaton), GitHub repository: https://github.com/XC-Li/Parallel_CellularAutomaton_Wildfire?tab=readme-ov-file (accessed 28.03.25).
[27] Historical Wildfire Data: 2006 to 2024, Government of Alberta, Canada, February. 2025, https://open.canada.ca/data/en/dataset/a221e7a0-4f46-4be7-9c5a-e29de9a3447e (accessed 02.04.25).
[28] A. B. Nordvik, Time window-of-opportunity strategies for oil spill planning and response, Pure Appl. Chem., Vol 71, No. 1, pp. 5-16, Great Britain, 1999.
[29] M. Gług, J. Wąs, Modeling of oil spill spreading disasters using combination of Langrangian discrete particle algorithm with Cellular Automata approach, Ocean Engineering, Volume 156, 2018, Pages 396-405.
[30] S. Naz, M. F. Iqbal, I. Mahmood, M. Allam, Marine oil spill detection using Synthetic Aperture Radar over Indian Ocean, Marine Pollution Bulletin, Volume 162, 2021, 111921.
[31] B. N. Bradford, P. J. Sanchez-Reyes, Automated oil spill detection with multispectral imagery, Proc. SPIE 8030, Ocean Sensing and Monitoring III, 80300L, 4 May 2011.
[32] L. Diana, J. Xu, L. Fanucci, Oil Spill Identification from SAR Images for Low Power Embedded Systems Using CNN. Remote Sens. 2021, 13, 3606.
[33] AICRAFT will deliver edge computing modules for South Korea's project for methane monitoring, AICRAFT Pty Ltd., Australia, May 28 2024, https://aicraft.com.au/news/aicraft-will-deliver-edge-computing-modules-for-south-koreas-project-for-methane-monitoring/ (accessed 02.04.25).
[34] B. Nolan, Illuminating the Invisible: How Ubotica is Revolutionising Dark Vessel Detection, Ubotica Technologies Limited, Dublin, Ireland, March 13 2025, https://ubotica.com/illuminating-the-invisible-how-ubotica-is-revolutionising-dark-vessel-detection/ (accessed 30.03.25).