

On-orbit change detection for remote sensing applications: Opportunities and challenges

Su Ye¹, Yuchen Tao¹, Xiucheng Yang^{2,3}

¹Zhejiang Key Laboratory of Agricultural Remote Sensing and Information Technology, Zhejiang University, Hangzhou 310058, China, ²Department of Biology, University of Victoria, Victoria, BC, Canada, ³Department of Geography, McGill University, Montreal, Quebec, Canada

Corresponding authors: Su Ye, su.ye@zju.edu.cn; and Xiucheng Yang, xiuchengyang@uvic.ca

Change detection is a fundamental task in remote sensing, supporting applications such as disaster response, land monitoring, infrastructure assessment, and security-related observation. Conventional change detection workflows generally depend on downlinking satellite imagery to ground stations, where preprocessing, analysis, interpretation, and product generation are performed. Although this ground-based paradigm has enabled reliable large-scale monitoring, it can introduce substantial latency between image acquisition and actionable information. This delay limits the value of satellite observations in time-critical scenarios where early alerts are essential. On-orbit change detection provides a promising alternative by enabling satellites to identify, prioritize, and report relevant surface changes directly onboard. This study reviews the opportunities and challenges of on-orbit change detection for remote sensing applications. We first examine the latency components of ground-based change detection, including observation lag, data lag, sensitivity lag, and production lag, and discuss how onboard processing may reduce or reshape these delays. We then synthesize traditional change detection methods, including image differencing, post-classification comparison, and time-series analysis, and evaluate how they can be adapted for resource-constrained spacecraft environments. Finally, we discuss key challenges for operational deployment, including lightweight processing, model updating, validation under lag-accuracy trade-offs, and change-category identification. The review shows that on-orbit change detection is unlikely to replace ground-based analysis entirely, but can provide earlier alerts, reduce unnecessary data downlink, and support more responsive Earth observation systems through staged and adaptive monitoring workflows.

Keywords – Change detection, near-real-time, on-orbit, onboard processing, satellite image analysis

1. INTRODUCTION

Change detection has long been a central task in remote sensing, enabling the systematic identification of temporal changes on Earth's surface from multi-temporal observations [1]-[3]. It supports a wide range of applications, including natural disaster assessment [4], [5], land-use monitoring [6], infrastructure development [7], and security-relevant event detection [8]. As satellite constellations continue to advance in spatial resolution, revisit frequency, and spectral capability, change detection systems must increasingly process large volumes of high-frequency, multi-temporal data. While this data enables more detailed and timely monitoring of dynamic Earth systems, they also introduce challenges related to data transmission, onboard and ground-based processing, and rapid decision-making [9].

Traditional ground-based change detection typically relies on transmitting acquired imagery to ground stations, where satellite data is preprocessed, analyzed, and interpreted. Although this workflow has supported a broad range of scientific, commercial, and operational applications, it can introduce significant latency between image acquisition and actionable insight. In time-critical scenarios, such as wildfire monitoring, flood assessment, military surveillance, disaster response, and rapid infrastructure damage evaluation, delays caused by downlink limitations and centralized processing can reduce the operational value of satellite observations [4], [10].

On-orbit change detection offers a promising alternative by enabling spacecraft to identify, prioritize, and potentially respond to relevant changes directly onboard. By performing preliminary analysis in space, on-orbit processing can reduce the downlink of redundant or low-priority data, optimize communication bandwidth, and accelerate the delivery of critical information to end users [11], [12]. This capability is particularly important for autonomous

satellite systems and distributed constellations, where rapid onboard decision-making can support adaptive sensing, event-driven imaging, and real-time mission planning [13], [14].

Despite its potential, on-orbit change detection remains technically challenging. Space-borne platforms typically operate under strict constraints in computational power, memory, energy consumption, and radiation tolerance [15], [16]. At the same time, change detection algorithms must remain robust to variations in illumination, viewing geometry, atmospheric conditions, seasonal effects, and sensor noise [1], [17], [5]. These requirements call for solutions that are not only accurate but also computationally-efficient and suitable for deployment on resource-constrained onboard hardware.

This study evaluates change detection algorithms for potential on-orbit deployment using satellite imagery. First, we synthesize existing change detection methods and the latency components of conventional ground-based processing workflows. We then review current adaptations of these methods for on-orbit change detection and assess their suitability for onboard implementation. Finally, we discuss key technical challenges and operational considerations associated with on-orbit change detection. By examining the requirements, limitations, and potential solutions for onboard processing, this work aims to identify methods capable of detecting meaningful surface changes with reduced latency and improved operational efficiency.

2. GROUND-BASED CHANGE DETECTION

Ground-based change detection latency can be understood as the cumulative delay between the actual occurrence of a surface change and the delivery of an operational change

alert. Adapted from the latency framework proposed by [4], this delay can be divided into four primary components when conventional ground-based processing is used for operational monitoring (Table 1). First, observation lag refers to the time required for a satellite sensor to acquire the first usable post-change observation, which is influenced by revisit frequency, cloud cover, atmospheric conditions, and viewing constraints. Second, data lag is the delay between image acquisition and the availability of analysis-ready products, including downlink, data transfer, preprocessing, calibration, atmospheric correction, and the generation of standard reflectance products. Third, sensitivity lag describes the time between the availability of the first usable post-change observation and the point at which the algorithm first detects the surface change. Finally, production lag is introduced by operational workflows such as map generation, product update schedules, quality control, database ingestion, and dissemination to end users. Together, these components indicate that latency in conventional ground-based change detection is not caused by a single bottleneck but by the combined effects of satellite observation timing, data availability, algorithmic sensitivity, and centralized production workflows. Depending on the sensor, algorithm, and operational configuration, total detection latency can range from five days [4] to several weeks or approximately two months in other Near-Real-Time (NRT) monitoring systems [18], [19]. It is also worth noting that many operational remote-sensing systems require multiple anomalous observations before confirming a change, in order to reduce false alarms and increase detection confidence [20], [21]. Although this confirmation requirement improves reliability, it can substantially increase overall detection latency, particularly when clear post-change observations are infrequent.

Table 1 – Latency components in ground-based change detection using the standard Harmonized Landsat and Sentinel-2 (HLS) dataset (30-meter resolution, 2-3 revisiting cycles, released by NASA) as an example, and the potential role of on-orbit processing.

Application	Target alert latency	Main change signal	Accuracy priority	Suitable onboard strategy	Key limitations
Wildfire	Minutes–hours	Thermal anomalies, smoke, spectral change	High initial recall	Thermal/event classifier + rapid alert	False alarms, cloud/smoke interference, revisit frequency
Flood	Hours	Rapid water extent increase	Early localization and extent estimation	Water classification + ROI downlink	Cloud obstruction for optical sensors, transient water signals
Earthquake damage	Minutes–hours	Structural/landscape anomalies	Early localization	Damage/anomaly detection + ROI downlink	Unseen damage signatures
Deforestation	Days	Persistent vegetation loss	High precision and lower false-alert rate	Time series analysis with repeated observations	Seasonal variability, selective logging, cloud contamination
Infrastructure change	Days–weeks	Persistent geometric/semantic change	High categorical accuracy	Infrastructure classification/comparison	Small-object detection and spatial-resolution requirements

The latency framework clarifies where onboard processing can provide benefit, but it does not determine which algorithms can realistically deliver that benefit. Data and production lag may be reduced only if change detection can be executed within spacecraft resource constraints. The following discussion therefore first summarizes the methodological principles of major change-detection approaches before evaluating their suitability for onboard implementation.




Figure 1 – Conceptual workflow of latency sources in ground-based and on-orbit change detection. Compared to four-component latency for ground-based processing, on-orbit change detection was only impacted by the revisiting cycle of satellites (observation lag) and algorithm acuteness (sensitivity lag).

Traditional change detection methods in remote sensing can generally be grouped into three categories: image differencing, post-classification comparison, and time-series analysis [1], [22]. These approaches differ in the dataset they require and the way they detect changes.

Image differencing is one of the simplest and most widely used change detection approaches. It compares two co-registered images acquired over the same area at different times by subtracting pixel values or spectral indices from one date to the other. For example, changes in vegetation, water, or burned area can be detected by differencing indices such as the Normalized Differencing Vegetation Index (NDVI) or Normalized Burning Ratio (NBR). For single-band processing, pixels with difference values exceeding a predefined or adaptive threshold are classified as changed, while the remaining pixels are treated as unchanged; for multiband inputs, each pixel is represented as a spectral vector at each acquisition time (also known as change vector analysis [23]), and the difference between the two vectors is used to describe both the magnitude and direction of change. The magnitude indicates the overall intensity of change, while the direction can provide information about the type or nature of change, such as vegetation loss, urban expansion, flooding, or burned area. The main advantage of image differencing is its computational simplicity, which makes it attractive for rapid processing and potential onboard implementation. However, it is highly sensitive to radiometric inconsistencies, illumination differences, atmospheric effects, seasonal variation, and image registration errors [22], [24]. Therefore, accurate preprocessing and threshold selection are critical for reliable performance.

Post-classification comparison detects change by first classifying each image independently into thematic land cover or land use classes using traditional machine learning or deep learning, and then comparing the resulting classification maps [25]. For example, an area classified as "forest" at the first date and "urban" at the second date would be labeled as a forest-to-urban transition. A major strength of this approach is that it provides not only the location of change but also the type of change, making it useful for applications such as land-use monitoring, deforestation assessment, agricultural mapping, and urban expansion analysis. It is also less directly affected by radiometric differences between images because change is measured at the class-label level rather than the raw pixel-value level. However, its accuracy depends heavily on the quality of the classification results at each time point. Classification errors from either date can propagate into the final change map, sometimes producing false changes. In addition, post-classification comparison usually requires reliable training data or robust classifiers, which may limit its suitability for fully autonomous onboard deployment unless pre-trained models or simplified classification schemes are used [26].

Time-series analysis uses multiple observations acquired over an extended period, rather than relying only on a bi-temporal image pair. This approach models the temporal behavior of pixels or regions and identifies change as deviations, breaks, trends, or anomalies in the observed time series [27]–[29]. Long-term satellite records are particularly useful for distinguishing abrupt disturbances, such as fire, flooding, and deforestation, from gradual trends or normal seasonal dynamics through the detection of temporal breaks (Fig. 2) [30]. Time series methods can also reduce false detections caused by temporary noise, cloud contamination, atmospheric effects, and seasonal variation by evaluating new observations within the context of historical surface behavior [31]. In this sense, time-series analysis frames change detection as a continuous monitoring problem, where each incoming observation is compared against an expected temporal pattern. However, time series approaches generally require larger data volumes, greater memory capacity, and higher computational resources than simple bi-temporal methods, which creates challenges for onboard implementation. For example, the Continuous Change Detection and Classification (CCDC) algorithm relies on historical observations to estimate and update temporal model coefficients, making it computationally demanding for resource-constrained platforms [30]. Stochastic Continuous Change Detection (S-CCD) reduces this burden by using recursive updating within a state-space modeling framework, lowering data storage and memory requirements to approximately one-tenth of those required by CCDC [19]. Nevertheless, S-CCD still requires local storage of model parameters and covariance matrices, which may remain substantial for full-scene or large-area onboard processing. Therefore, although time-series analysis offers strong robustness for operational change monitoring, its storage, memory, and computational requirements make full onboard deployment challenging without further simplification, compression, or lightweight model design.




Figure 2 – The conceptual figure for detecting change based upon detecting consecutive anomalies using a peek window for a satellite time-series dataset, which was adopted in the popular algorithms such as CCDC and BFAST.

3. ON-ORBIT CHANGE DETECTION

3.1 Algorithm descriptions

The preceding section establishes both the latency components of conventional ground-based monitoring and the methodological principles of the three major change-detection families. These methods, however, were largely developed under ground-processing assumptions and therefore cannot be transferred directly to spacecraft. Table 2 compares representative onboard change detection approaches, including adaptations of three traditional change-detection method families and one emerging approach based on event-specific detection. Note that the qualitative ratings in Table 2 indicate relative onboard resource requirements rather than standardized benchmark results. Actual computational and memory demands depend on sensor resolution, scene size, preprocessing requirements, model architecture, hardware platform, and implementation. Potential downlink reduction additionally depends on whether onboard outputs replace full-image transmission or are transmitted alongside the original imagery.

For image differencing and change vector analysis, onboard adaptation typically shifts the comparison from raw pixel values to compact feature, score, or semantic representations. Rather than storing and comparing full multi-band images, these methods use lightweight models to encode each observation into a lower-dimensional representation or to assign each pixel or image tile a change likelihood score. Change is then detected by measuring the distance between

feature vectors, thresholding change scores, or selecting pixels with a high probability of change. This preserves the basic logic of bi-temporal differencing while reducing memory, computation, and downlink requirements, making the approach more suitable for resource-constrained onboard implementation. For example, Bui et al. [32] used Cloud-Net for cloud removal and Change-Net to score multi-spectral pixels according to their likelihood of change, so that only changed or high-priority pixels are encoded and transmitted to the ground. In this way, changed regions can be prioritized for downlink instead of transmitting entire scenes. The computational simplicity of image differencing makes it attractive for onboard deployment, but its apparent efficiency can be misleading if reliable operation requires expensive radiometric normalization, geometric co-registration, cloud screening, or multi-observation confirmation. Thus, the onboard cost should be evaluated for the complete processing chain rather than the differencing operator alone.

For post-classification comparison, onboard change detection is often implemented through event-specific classification. In this case, the spacecraft first identifies a target class or event, such as flood water, volcanic activity, snow, ice, or cloud cover, and then uses the classification result to support autonomous response or selective transmission. A representative example is the Autonomous Spacecraft Experiment (ASE) on EO-1, in which an onboard thermal classifier analyzed Hyperion hyperspectral data to detect volcanic thermal anomalies and autonomously support follow-up observations of active volcanism [34]. This approach is closely related to traditional post-classification comparison because change is inferred from semantic class labels or event detections rather than direct pixel differencing alone. Its main advantage is interpretability and operational relevance, as detected classes or events can directly trigger mission actions such as re-tasking, prioritization, or selective downlink. However, its limitation is that it usually detects only predefined event types for which classifiers or rules have been designed, and they may not generalize well to unexpected or ambiguous changes. Semantic classifiers provide operationally interpretable outputs, but their advantage is strongest when the target event is known in advance. Increasing the number of change classes increases both model complexity and the risk of domain shift; consequently, a high-performing ground-based multi-class model may be less suitable onboard than a smaller event-specific classifier.

Table 2 – Comparison of on-orbit change detection approaches for onboard implementation.

Approaches	Typical onboard model representation	Compute demand	Memory demand	Potential downlink reduction	Reference examples
Image differencing	Previous image/index or compact reference	Low	Low–moderate	High only when ROIs/alerts are sent	[32], [33]
Post-classification comparison	Class labels/probabilities	Moderate	Moderate	High if semantic alerts replace imagery	[34], [35]
Time-series analysis	Model state/covariance/history summary	Moderate	High	Moderate	[36], [37], [38]
Event detector	Model weights	Moderate	Low–moderate	High when detections/ROIs replace imagery	[39], [12]

For time-series analysis, onboard methods compare new observations against a historical representation of the same location. Rather than relying on a single pre-event and post-event image pair, these approaches use previous observations to construct a priori model of baseline conditions and then determine whether the current observation is anomalous or novel. For example, unsupervised methods such as [36] encode image tiles into a latent feature space and assign novelty scores according to their distance from previous observations. Similarly, Xu et al. [38] proposed a general on-orbit Earth surface anomaly detection framework that constructs lightweight historical priors using large vision model features and performs efficient onboard anomaly prediction through a dictionary lookup-style comparison. These methods make onboard change detection especially useful for identifying unexpected or diverse event types, because they can flag unusual changes without requiring predefined change categories. However, this approach still requires careful management of onboard storage, temporal baselines, and false alarms caused by seasonal variation, sensor noise, or imaging differences. In addition, because the historical prior is constructed at the image or patch-feature level rather than as a dense pixel-level time series, it may reduce pixel-level sensitivity and localization precision for subtle or spatially heterogeneous changes. Time-series algorithms offer greater robustness to normal temporal variability, but their historical context creates the very memory and update requirements that conflict with onboard deployment. Recursive models partially resolve this contradiction, suggesting that the key design variable is not simply algorithm family but the form in which historical information is retained.

In addition to adaptations of conventional change-detection methods, a complementary onboard strategy is the use of lightweight event detectors, which directly identify predefined, time-critical events such as wildfires, floods, or thermal anomalies from newly acquired observations. Unlike general change-detection methods that explicitly compare multiple dates, event detectors can infer the presence of an event from a single observation using compact CNNs or other lightweight classifiers, thereby avoiding the storage and processing of historical imagery. This makes them particularly attractive when rapid alerts are more important than comprehensive change characterization. For example, lightweight CNNs have been developed for onboard real-time classification of thermal hotspots from minimally preprocessed satellite data, demonstrating the feasibility of rapid event screening under spacecraft resource constraints [39][12]. However, these approaches are generally optimized for predefined targets and therefore have limited ability to detect unexpected or previously unseen change types. Consequently, event detectors are best viewed as an application-specific complement to general bi-temporal or time-series change detection, particularly for latency-critical applications.

3.2 Application-dependent deployment requirements

The suitability of an on-orbit change detection approach depends strongly on the operational requirements of the target application, particularly the acceptable detection latency, false-alert tolerance, spatial detail, and level of semantic information required (Table 3). Time-critical events such as wildfires, floods, and earthquake damage generally place greater emphasis on early detection and localization because delays can substantially reduce the operational value of satellite observations for emergency response. For wildfire monitoring, thermal anomalies, smoke, and rapid spectral changes can provide early evidence of an event, favoring high initial recall and rapid onboard event detection, although false alarms, atmospheric interference, and satellite revisit frequency remain important limitations. Flood monitoring similarly benefits from rapid identification of increases in water extent, for which onboard water classification followed by selective downlink of Regions of Interest (ROIs) can provide timely information without requiring transmission of complete scenes. Earthquake damage presents a different challenge because structural and landscape changes can be spatially heterogeneous and may differ substantially from previously observed damage patterns. Consequently, onboard processing for such events may be most useful for early localization of potentially affected areas and prioritization of ROIs for subsequent ground-based analysis rather than comprehensive damage assessment onboard.

In contrast, deforestation and infrastructure change generally evolve over longer operational timescales and therefore place greater emphasis on detection confidence and semantic accuracy than on issuing an alert from the earliest possible observation (Table 3). Deforestation monitoring can benefit from repeated observations and time-series analysis to distinguish persistent vegetation loss from seasonal variability or other temporary changes, although maintaining historical information increases onboard memory and computational requirements. Infrastructure monitoring may similarly tolerate delays of days to weeks but requires sufficient spatial resolution and categorical accuracy to identify persistent geometric or semantic changes, particularly for small objects. These contrasts demonstrate that there is no universally optimal strategy for on-orbit change detection. Instead, method selection represents an application-dependent trade-off among latency, accuracy, computational and memory requirements, and the information content of the onboard product. For rapid-onset disasters, lightweight detection and selective ROI downlink may provide greater operational value than generating a fully refined change product onboard, whereas persistent land change applications can justify additional temporal context and processing to reduce false alerts and improve confidence. Table 3 summarizes these differences and highlights how the timescale and characteristics of the target change should guide the design of the onboard detection and reporting strategy.

Table 3 – Application-dependent requirements and onboard strategies for on-orbit change detection.

Application	Target alert latency	Main change signal	Accuracy priority	Suitable onboard strategy	Key limitations
Wildfire	Minutes–hours	Thermal anomalies, smoke, spectral change	High initial recall	Thermal/event classifier + rapid alert	False alarms, cloud/smoke interference, revisit frequency
Flood	Hours	Rapid water extent increase	Early localization and extent estimation	Water classification + ROI downlink	Cloud obstruction for optical sensors, transient water signals
Earthquake damage	Minutes–hours	Structural/landscape anomalies	Early localization	Damage/anomaly detection + ROI downlink	Unseen damage signatures
Deforestation	Days	Persistent vegetation loss	High precision and lower false-alert rate	Time series analysis with repeated observations	Seasonal variability, selective logging, cloud contamination
Infrastructure change	Days–weeks	Persistent geometric/semantic change	High categorical accuracy	Infrastructure classification/comparison	Small-object detection and spatial-resolution requirements

4. CURRENT CHALLENGES

4.1 Lightweight processing under onboard resource constraints

Lightweight processing is a central challenge for on-orbit change detection because spacecraft operate under strict size, weight, power, memory, thermal, radiation-tolerance, and communication constraints. Unlike ground-based workflows, which can rely on large storage systems and high-performance processors, onboard systems must execute reliably within limited resources and predictable time budgets. This is particularly challenging for multi-temporal change detection, which may require storing previous observations, maintaining reference baselines, comparing new images with historical data, and filtering false alarms caused by illumination, atmospheric effects, seasonal variation, or sensor noise. As a result, onboard systems often need simplified models, reduced input dimensionality, efficient feature extraction, tile-based processing, or compressed representations, rather than computationally expensive full-scene analysis.

Onboard hardware now ranges from radiation-hardened CPUs and FPGAs to heterogeneous CPU–GPU–VPU platforms and increasingly specialized AI accelerators [15]. Traditional FPGA-based systems provide reliable and energy-efficient parallel processing, while newer heterogeneous architectures enable more sophisticated machine-learning inference. Recent missions such as Φ sat-1 and Φ sat-2 demonstrate this transition by performing AI-based cloud screening and other Earth-observation tasks directly onboard [40][41], while platforms such as Unibap iX5 illustrate the growing use of combined CPU, GPU, and VPU resources [42]. However, greater computational capability generally increases power consumption, thermal-

management requirements, and the need for fault-mitigation mechanisms. Radiation-induced errors also require safeguards such as error correction, redundancy, watchdog mechanisms, or configuration scrubbing, making hardware and algorithm selection closely coupled.

Memory and data movement are equally important. Bi-temporal methods may require stored reference observations, while time-series approaches retain historical states or model coefficients. Compact representations, tile-based processing, recursive updating, and low-precision inference can reduce these requirements. For example, S-CCD reduces storage and memory demand to approximately one-tenth of CCDC by recursively updating temporal-model states rather than retaining a more extensive historical representation [19][43]. Besides, preprocessing costs must also be considered: image registration, radiometric correction, cloud screening, and normalization may substantially increase the end-to-end burden even when the change algorithm itself is simple.

Communication limitations provide a major motivation for onboard processing. Rather than producing fully refined products in orbit, practical systems can transmit only regions of interest, anomaly scores, thumbnails, or high-priority disturbance alerts. This favors staged processing in which lightweight onboard screening first identifies potentially relevant observations, followed by more detailed ground-based analysis. Consequently, methods should be assessed jointly in terms of accuracy, latency, memory, power, preprocessing demand, radiation resilience, and expected downlink reduction. The most suitable onboard method is therefore not necessarily the most accurate standalone algorithm but the one that provides sufficient and reliable information within the mission-specific resource envelope.

4.2 Model updates

Model updating is another important challenge for on-orbit change detection because the conditions under which a model operates can change over time. A model trained on historical ground-based datasets may not generalize well to new geographic regions, seasons, illumination conditions, sensor degradation, atmospheric effects, land cover types, or disturbance patterns. This is especially problematic for change detection because the boundary between true change and normal variability is often context-dependent. For example, seasonal vegetation cycles, agricultural rotation, snow cover, drought stress, or differences in viewing geometry may appear as change if the onboard model is not adapted to local or temporal conditions. Therefore, an onboard change detection model must either be robust enough to handle a wide range of operating conditions or be periodically updated to reflect new observations and changing environments.

However, updating models onboard is difficult because satellites have limited computational resources, restricted memory, intermittent communication windows, and strict reliability requirements. Full retraining is usually impractical in orbit, especially for deep learning models that require large labeled datasets and high-performance hardware. A more feasible strategy is to train and validate models on the ground, then uplink lightweight model updates, calibrated thresholds, new class prototypes, or region-specific parameters to the satellite. Incremental learning, domain adaptation, self-supervised feature updating, and uncertainty-based monitoring may also help onboard systems adapt without complete retraining. Still, model updates must be carefully controlled because an unreliable update could degrade detection performance across an entire mission. As a result, on-orbit change detection systems need robust update mechanisms, including version control, validation before deployment, rollback capability, and safeguards against model drift.

4.3 Validation

Validation is a major challenge for on-orbit change detection because performance cannot be assessed solely using final map accuracy. In operational monitoring, the value of a change detection system depends not only on how accurately it identifies change, but also on how quickly it delivers reliable information. A method that produces a highly accurate change map several weeks after an event may be less useful for time-critical applications than a less complete but earlier alert. Therefore, validation should account for the relationship between latency and accuracy rather than relying on a single static accuracy score. Metrics such as omission error, commission error, precision, recall, and F1 score can be evaluated as functions of latency to show how detection performance evolves from the first observable change signal to a stable performance plateau. This lag–accuracy curve, as discussed by [44], is especially important for on-orbit processing because it captures the temporal performance of the entire monitoring workflow, including observation, onboard analysis, communication, and product delivery (Fig. 3).




Figure 3 – The conceptual figure for a lag–accuracy curve. Following [44], the first-alert lag is assigned to the date when 2% of all change pixels are detected, indicating the earliest time at which an algorithm can begin issuing land change alerts. The level-off lag represents the time when the F1 score stabilizes. It is calculated from the fitted sigmoid curve as the date at which the rotated curve reaches its maximum y-value after a 45° rotation around the line connecting the curve's beginning and end points.

However, constructing reliable lag–accuracy curves for on-orbit change detection is difficult. It requires accurate reference information on both the location and timing of true changes, ideally from higher-resolution imagery, dense time series, field reports, or authoritative event records. If the reference change timing is uncertain, the estimated latency of the onboard system may be biased. Validation is also complicated by the trade-off between early detection and false alarms: reducing latency often increases commission errors, while requiring stronger confirmation can improve reliability but delay alerts. For onboard systems, this trade-off is particularly important because false positives may waste limited downlink bandwidth or trigger unnecessary follow-up imaging, whereas false negatives may cause critical events to be missed. As a result, validation frameworks for on-orbit change detection should report not only final accuracy but also initial detection delay, time to stable performance, false-alert burden, and the operational usefulness of early alerts under different latency thresholds.

The preceding discussion treats change detection primarily as a binary decision: whether sufficient evidence exists to issue an alert. Operational systems, however, frequently require a second decision after detection, regarding what type of change has occurred. These two tasks need not reach useful accuracy at the same time. A system may detect an anomalous surface transition from the first post-event observation while requiring additional spectral, temporal, or contextual evidence before reliably assigning a semantic category. This distinction motivates the following challenge of change-category identification.

4.4 From change detection to change-category identification

Identifying the category of change is a major challenge for on-orbit change detection because early alerts often contain limited evidence about what type of event has occurred. Detecting that a surface has changed is generally easier than determining whether the change is caused by wildfire,

flooding, deforestation, construction, crop rotation, storm damage, seasonal vegetation dynamics, or sensor-related noise. This distinction becomes even more difficult under low-latency requirements. At the beginning of the lag–accuracy curve, only one or a few post-change observations may be available, so the algorithm may detect an anomalous signal but lack enough temporal context to assign a reliable change category. For example, a sudden decrease in vegetation index could indicate fire damage, flooding, harvest, drought stress, insect disturbance, or cloud-shadow contamination depending on location, season, spectral response, and spatial pattern. As a result, early on-orbit products may be better suited for generating generic "change" or "anomaly" alerts, while more specific change-category labels may require additional observations, contextual information, or ground-based refinement.

From a lag–accuracy perspective, change-category identification introduces a second trade-off beyond simple change/no-change detection. An onboard system may

achieve short detection lag by issuing broad change alerts, but category-level accuracy may remain low until more evidence is accumulated. As latency increases, additional observations can improve discrimination among change types by revealing temporal persistence, recovery behavior, spectral-change direction, object shape, spatial extent, and contextual relationships with roads, rivers, vegetation, or built-up areas. However, waiting for these features improves classification confidence at the cost of reduced timeliness. A useful example is the ground-based platform developed by [21], which successfully identifies forest disturbance types in an operational setting but still has an approximately 10-day latency (Fig. 4). A promising solution is a staged reporting framework: first, transmit low-latency generic alerts or regions of interest; second, update these alerts with probabilistic change-category labels as more observations become available; and finally, refine the category and severity on the ground using higher-resolution data, ancillary information, and more computationally intensive models.




Figure 4 – Different forest disturbance types labeled by the NRT platform "China's Forest Disturbance Watch" (<https://forestwatch.com.cn/>).

5. CONCLUSION

On-orbit change detection offers a promising pathway toward faster and more autonomous remote sensing by shifting part of the analysis from ground-based systems to the spacecraft itself. This review shows that onboard processing is especially useful for reducing data lag and production lag by allowing satellites to identify, prioritize, and transmit only high-value observations, alerts, metadata, or selected regions of interest instead of downlinking all acquired imagery. However, it cannot fully remove observation lag, which is governed by satellite revisit frequency, cloud cover, and sensor availability, and sensitivity lag remains dependent on algorithm performance and the detectability of early change signals. Traditional change detection methods, including image differencing, post-classification comparison, and time-series analysis, provide useful foundations for on-orbit implementation, but they must be adapted to strict constraints in computation, memory, energy, communication, and reliability. Future operational systems will require lightweight algorithms, adaptive model updating, latency-aware validation, and staged reporting workflows that combine rapid onboard alerts with more detailed ground-based processing. Overall, on-orbit change detection should be viewed as a complement to conventional ground-based processing, enabling earlier awareness, more efficient data transmission, and more responsive monitoring systems.

ACKNOWLEDGEMENT

The authors are grateful for the financial support provided by National Natural Science Foundation of China (42401393).

REFERENCES

- [1] D. Lu, P. Mausel, E. Brondizio, and E. Moran, "Change detection techniques," *Int. J. Remote Sens.*, vol. 25, no. 12, pp. 2365-2401, 2004.
- [2] C. E. Woodcock, T. R. Loveland, M. Herold, and M. E. Bauer, "Transitioning from change detection to monitoring with remote sensing: A paradigm shift," *Remote Sens. Environ.*, vol. 238, p. 111558, 2020.
- [3] R. E. Kennedy et al., "Remote sensing change detection tools for natural resource managers: Understanding concepts and tradeoffs in the design of landscape monitoring projects," *Remote Sens. Environ.*, vol. 113, no. 7, pp. 1382-1396, 2009.
- [4] S. Ye, Z. Zhu, and J. W. Suh, "Leveraging past information and machine learning to accelerate land disturbance monitoring," *Remote Sens. Environ.*, vol. 305, p. 114071, 2024.

- [5] S. Ye et al., "Detecting subtle change from dense Landsat time series: Case studies of mountain pine beetle and spruce beetle disturbance," *Remote Sensing of Environment*, vol. 263, 2021. doi: 10.1016/j.rse.2021.112560.
- [6] A. H. Chughtai, H. Abbasi, and I. R. Karas, "A review on change detection method and accuracy assessment for land use land cover," *Remote Sens. Appl. Soc. Environ.*, vol. 22, p. 100482, 2021.
- [7] J. W. Suh, Z. Zhu, and Y. Zhao, "Monitoring construction changes using dense satellite time series and deep learning," *Remote Sens. Environ.*, vol. 309, p. 114207, 2024.
- [8] M. J. Carlotto and S. Member, "Detection and Analysis of Change in Remotely Sensed Imagery with Application to Wide Area Surveillance," vol. 6, no. 1, pp. 189-202, 1997.
- [9] X. Liu et al., "On-orbit image processing technology for intelligent remote sensing satellites : Progress , challenges , and opportunities," *Aerosp. Sci. Technol.*, vol. 174, p. 111859, 2026, doi: 10.1016/j.ast.2026.111859.
- [10] M. E. Brown, M. L. Carroll, and V. M. Escobar, "User needs and assessing the impact of low latency NASA Earth observation data availability on societal benefit," *Space Policy*, vol. 30, no. 3, pp. 135-137, 2014.
- [11] R. Del Prete et al., "Optimizing deep learning models for on-orbit deployment through neural architecture search," pp. 1-13, 2025.
- [12] F. Ip et al., "Flood detection and monitoring with the Autonomous Sciencecraft Experiment onboard EO-1," *Remote Sens. Environ.*, vol. 101, no. 4, pp. 463-81, 2006.
- [13] B. Cichy, G. Rabideau, D. Tran, and R. Sherwood, "Monitoring active volcanism with the Autonomous Sciencecraft Experiment on EO-1," vol. 101, pp. 427-446, 2006, doi: 10.1016/j.rse.2005.08.007.
- [14] S. Chien et al., "The EO-1 autonomous science agent," in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 2004, pp. 420-427.
- [15] L. Diana and P. Dini, "Review on Hardware Devices and Software Techniques Enabling Neural Network Inference Onboard Satellites," pp. 1-29, 2024.
- [16] A. Duggan, B. Andrade, and H. Afli, "Advancing Earth observation : a survey on AI- powered image processing in satellites," vol. 7254, 2025, doi: 10.1080/22797254.2025.2567921.
- [17] Z. Zhu, "Change detection using landsat time series: A review of frequencies, preprocessing, algorithms, and applications," *ISPRS J. Photogramm. Remote Sens.*, vol. 130, pp. 370-384, 2017.
- [18] R. Shang et al., "Near-real-time monitoring of land disturbance with harmonized Landsats 7-8 and Sentinel-2 data," *Remote Sens. Environ.*, vol. 278, p. 113073, 2022.
- [19] S. Ye, J. Rogan, Z. Zhu, and J. R. Eastman, "A near-real-time approach for monitoring forest disturbance using Landsat time series: stochastic continuous change detection," *Remote Sensing of Environment*, vol. 252, 2021. doi: 10.1016/j.rse.2020.112167.
- [20] A. H. Pickens et al., "Rapid monitoring of global land change," *Nat. Commun.*, vol. 16, no. 1, p. 8948, 2025, doi: 10.1038/s41467-025-64014-9.
- [21] R. Liao et al., "Remote Sensing of Environment Near-real-time mapping for causal agents of forest disturbances in China using harmonized Landsat and Sentinel-2 dataset," vol. 342, no. May, 2026, doi: 10.1016/j.rse.2026.115461.
- [22] M. Hussain, D. Chen, A. Cheng, H. Wei, and D. Stanley, "Change detection from remotely sensed images: From pixel-based to object-based approaches," *ISPRS J. Photogramm. Remote Sens.*, vol. 80, pp. 91-106, 2013, doi: 10.1016/j.isprsjprs.2013.03.006.
- [23] J. Chen, X. Chen, X. Cui, and J. Chen, "Change vector analysis in posterior probability space: A new method for land cover change detection," *IEEE Geosci. Remote Sens. Lett.*, vol. 8, no. 2, pp. 317-321, 2010.
- [24] S. Ye, D. Chen, and J. Yu, "ISPRS Journal of Photogrammetry and Remote Sensing A targeted change-detection procedure by combining change vector analysis and post-classification approach," vol. 114, pp. 115-124, 2016, doi: 10.1016/j.isprsjprs.2016.01.018.
- [25] J. Rogan and D. Chen, "Remote sensing technology for mapping and monitoring land-cover and land-use change," vol. 61, no. July 2003, pp. 301-325, 2004, doi: 10.1016/S0305-9006(03)00066-7.
- [26] M. Chen and W. Jiang, "SCDVit: Semantic Change Detection Based on Sam-Vit and Semantic Consistency," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 18, pp. 2269-2282, 2025, doi: 10.1109/JSTARS.2024.3514926.
- [27] Z. Zhu et al., "Continuous monitoring of land disturbance based on Landsat time series," *Remote Sens. Environ.*, 2019, doi: 10.1016/j.rse.2019.03.009.
- [28] Z. Zhu and C. E. Woodcock, "Continuous change detection and classification of land cover using all available Landsat data," *Remote Sens. Environ.*, vol. 144, pp. 152-171, 2014.
- [29] J. Verbesselt, R. Hyndman, G. Newnham, and D. Culvenor, "Detecting trend and seasonal changes in satellite image time series," *Remote Sens. Environ.*, vol. 114, no. 1, pp. 106-115, 2010.

- [30] Z. Zhu et al., "Continuous monitoring of land disturbance based on Landsat time series," *Remote Sens. Environ.*, vol. 238, p. 111116, 2020.
- [31] Z. Zhu, S. Qiu, and S. Ye, "Remote sensing of land change: A multifaceted perspective," *Remote Sens. Environ.*, vol. 282, p. 113266, 2022.
- [32] V. Bui, T. Q. Dinh, I. Leyva-mayorga, S. R. Pandey, E. Lagunas, and P. Popovski, "On-board Change Detection for Resource-efficient Earth Observation with LEO Satellites," *GLOBECOM 2023 – 2023 IEEE Glob. Commun. Conf.*, pp. 3971–3976, 2023, doi: 10.1109/GLOBECOM54140.2023.10437543.
- [33] G. Guerrisi, F. Del Frate, and G. Schiavon, "Satellite On-Board Change Detection via Auto-Associative Neural Networks," 2022.
- [34] A. G. Davies et al., "Monitoring active volcanism with the autonomous sciencecraft experiment on EO-1," *Remote Sens. Environ.*, vol. 101, no. 4, pp. 427–446, 2006.
- [35] T. Doggett et al., "Autonomous detection of cryospheric change with hyperion on-board Earth Observing-1," vol. 101, pp. 447–462, 2006, doi: 10.1016/j.rse.2005.11.014.
- [36] V. Růžicka et al., "RaVEn: unsupervised change detection of extreme events using ML on – board satellites," *Sci. Rep.*, no. 0123456789, pp. 1–10, 2022, doi: 10.1038/s41598-022-19437-5.
- [37] K. Tang, Q. Wang, F. Xu, Z. Gu, X. Chen, and J. Chen, "Unsupervised Detection of Disaster-Affected Areas by Combining the Strengths of Change Detection and Anomaly Detection: Target for On-Orbit Application," 2026, doi: 10.34133/remotesensing.0929.
- [38] J. Xu et al., "Toward a Novel Method for General On-Orbit Earth Surface Anomaly Detection Leveraging Large Vision Models and Lightweight Priors," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, pp. 1–21, 2024, doi: 10.1109/TGRS.2024.3432749.
- [39] G. Meoni et al., "E2E: Onboard satellite real-time classification of thermal hotspots events on optical raw data," *Astrodynamics*, vol. 9, no. 3, pp. 447–463, 2025.
- [40] G. Giuffrida et al., "CloudScout: A deep neural network for on-board cloud detection on hyperspectral images," *Remote Sens.*, vol. 12, no. 14, p. 2205, 2020.
- [41] N. Melega et al., "Development and implementation of the Φsat-2 mission," in *Small Satellites Systems and Services Symposium (4S 2024)*, SPIE, 2025, pp. 1244–1254.
- [42] S. Daniel, C. Umut, K. Daniel, and F. Eric, "Unibap iX5 platform with Xilinx Kintex FPGA front-end for payload sensors and sub systems integration," in *Small Satellites Systems and Services Symposium (4S 2024)*, SPIE, 2025, pp. 1894–1908.
- [43] S. Ye and Y. Hu, "Efficient continuous land change monitoring with pyxccd and S-CCD 2.0," 2026.
- [44] E. L. Bullock et al., "Timeliness in forest change monitoring: A new assessment framework demonstrated using Sentinel-1 and a continuous change detection algorithm," *Remote Sens. Environ.*, vol. 276, p. 113043, 2022.

AUTHORS



SU YE received his B.Sc. in forestry from Beijing Forestry University, China, M.Sc. in environmental studies from Queen's University, Canada, and Ph.D. in geography from Clark University, USA. He is currently an assistant professor at Zhejiang University, China. He has contributed to the development of multiple remote sensing algorithms and satellite-based monitoring and analytical platforms, including Mapping Africa, the CONUS Disturbance Watcher, and the Land Change Modeler in TerrSet. He is also the primary developer and maintainer of pyxccd, a Python package for satellite-based time series analysis. His research interests include remote sensing of land-system change, disturbance monitoring, and the development of scalable geospatial methods for environmental and land-resource applications.



YUCHEN TAO is currently a master's student at the College of Environmental and Resource Sciences, Zhejiang University, China. He received his B.Sc. in geographic information science from Nanjing Forestry University, China. His current research focuses on forest disturbance mapping and attribution, using dense satellite time series and the Continuous Change Detection and Classification (COLD) algorithm to detect, characterize, and attribute forest changes caused by both natural and anthropogenic drivers. His work aims to contribute to sustainable forest management and a better understanding of terrestrial ecosystem dynamics under environmental change.



XIUCHENG YANG received his Ph.D. in geoinformation from the University of Strasbourg in 2019. He was a postdoctoral researcher at the University of Connecticut, USA, from 2020 to 2025. He is currently a senior research fellow at the University of Victoria and McGill University in Canada. He has published more than 40 papers. His research focuses on satellite time series analysis, coastal wetlands, land cover change, ecosystem disturbance and resilience, and GeoAI.