drones

MDPI

Article

A UAV–Satellite Hybrid Pipeline for Wildfire Detection and Dynamic Perimeter Prediction

Hossein Keshmiri * and Khan A. Wahid

Electrical and Computer Engineering Department, University of Saskatchewan, Saskatoon, SK S7N 5A9, Canada; khan.wahid@usask.ca
* Correspondence: hossein.keshmiri@usask.ca

Highlights

What are the main findings?

- A hybrid pipeline successfully integrates multi-UAV aerial detection with satellite-based deep learning for comprehensive wildfire management.
- A specialized dataset for the Canadian boreal context enables the SE-ResUNet model to achieve high-quality dynamic perimeter prediction.

What are the implications of the main findings?

- The TinyML and Dual LoRa architecture permits persistent, low-latency surveillance in remote regions, reducing standard drone energy constraints.
- The validated simulation framework proves that accurate daily fire spread forecasting is achievable even in data-scarce environments, facilitating proactive emergency response.

Abstract

Effective wildfire management demands seamless integration of real-time detection and long-term spread forecasting. This paper proposes a novel power-efficient UAV–satellite hybrid pipeline that synergizes the agility of UAVs with the scale of satellite intelligence. The system begins with a dashboard-guided, multi-UAV detection module that scores fire likelihood from historical satellite data and enables scalable, energy-efficient deployment with low-latency onboard processing. This aerial component ensures persistent surveillance and reliable ignition detection, supported by a Dual LoRa (Long Range) communication scheme for robust and low-power connectivity. It achieves an F1-score of 97.4% while minimizing power consumption to extend operational flight times. Following detection, the pipeline transitions to a dynamic perimeter-prediction phase utilizing a custom Canadian boreal dataset. We employ a Squeeze-and-Excitation Residual U-Net (SE-ResUNet) to model spatiotemporal fire propagation based on static terrain and dynamic environmental features. The model was validated using a dynamic simulation framework that evaluates temporal consistency and convergence behavior against final cumulative burned-area masks, effectively addressing the absence of daily ground truth. Under these conditions, the model achieves a recall of 84% and an AUC of 0.97, demonstrating a strong capability to delineate active fire fronts. By coupling dashboard-driven UAV sensing with satellite-based predictive modeling, this work establishes a modular, foundational framework to support data-scarce forecasting in modern wildfire management.

Keywords: multi-UAV; wildfire detection; wildfire spread; dynamic prediction; hybrid pipeline

Check for updates

Academic Editors: Abraham Mejia-Aguilar, Santos Miguel Orozco Soto and Christian Wankmüller

Received: 30 January 2026
Revised: 31 March 2026
Accepted: 1 April 2026
Published: 4 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Wildfires are natural phenomena that historically play an important role in maintaining ecosystem balance. However, the adverse effects of climate change, largely driven by prolonged and environmentally destructive human activities, have disrupted this balance, causing wildfires to occur more frequently and with greater intensity and spread. In 2023, approximately 15 million hectares of land in Canada were burned by wildfires—more than double the previous national record [1]. Furthermore, during 2024–2025 in North America, about 46,000 km^2 and 36,000 km^2 burned in Canada and the United States, respectively, exceeding historical averages and reflecting elevated fire activity across the region [2]. As a result, wildfires have evolved from a natural ecological process into a major threat to both environmental sustainability and human life.

In recent decades, researchers have developed predictive methods to detect and control wildfires. Thangavel et al. [3] investigated space-borne artificial intelligence for rapid wildfire detection by implementing onboard autonomous processing on satellite-suitable hardware for disaster management systems. Tong et al. [4] presented a UAV-assisted monitoring system that employed convolutional neural networks and transformers to detect small fires from low-altitude imagery. Fan et al. [5] introduced an integrated framework for detecting both fire and smoke from UAV imagery using deep learning with mixed data augmentation. Their method improved accuracy and inference speed on UAV datasets by refining backbone features and incorporating shape-aware loss functions for irregular targets. Seydi et al. [6] proposed a deep learning framework trained on Landsat-8 imagery from Australia, Central Africa, Ukraine, and the Amazon rainforest, integrating optical and thermal modalities to improve active fire detection performance. In [7], the authors investigated active fire detection in Eastern Australia using support vector machines trained on multispectral satellite data to achieve robust classification of burned areas.

Based on the literature reviewed in Section 2, small battery-powered UAVs and satellite data are among the most widely used tools for wildfire detection, monitoring, and prediction. As summarized in Table 1, each platform exhibits distinct advantages and limitations. Owing to their high maneuverability, UAVs equipped with lightweight sensors can fly over target areas and acquire high-resolution, real-time information, such as imagery and environmental measurements, at user-defined distances. They can perform onboard processing or offload data to external servers for further analysis. Moreover, cooperative UAV networks can operate collaboratively to expand coverage, follow scheduled workflows, and achieve scalable, low-cost deployments. However, the benefits of flexible operation, high spatial resolution, and real-time acquisition come at the expense of limited onboard power and computational resources. These constraints restrict flight endurance and may reduce decision-making capability when UAVs are used independently for wildfire monitoring. Furthermore, because wildfire prediction requires long-term observations of geospatial variables across large regions, UAVs alone are generally insufficient for large-scale, long-duration prediction tasks.

Table 1. UAVs versus satellite data.

Desired Attributes	UAV	Satellite
Temporal adaptability	On-demand/high revisit	Fixed/low revisit
Data availability	Low latency (seconds to minutes)	High latency (hours to days)
Spatial resolution	Moderate to high	Low to moderate [1]
Scalability	High	Low
Implementation cost	Low	High
Operational endurance	Minutes to hours	Continuous
Ground coverage	Limited	Approximately all Earth

[1] Although several satellites such as Sentinel-2 (10–20 m) and Landsat-8/9 (30 m) provide comparatively high-resolution optical data, they suffer from reduced temporal coverage (e.g., Sentinel-2 available only since 2015; Landsat-8 since 2013), longer revisit cycles, and cloud limitations. Therefore, operational long-term wildfire applications predominantly rely on moderate-resolution sensors such as MODIS (250–500 m) and VIIRS (375 m) [8,9].

Satellite-driven data, on the other hand, is a valuable source of information that can be used for wildfire detection and prediction. Historical wildfire data, such as location, time, duration, and burned geometry, along with meteorological and geospatial data for approximately anywhere on Earth, can be acquired from various satellite databases to train intricate ML models serving the goal of wildfire detection as well as prediction. However, most standard satellite data products are available with a notable latency, typically ranging from 3 to 6 h relative to real-time [10,11]. Furthermore, these sensors often suffer from moderately coarse spatial resolutions, generally ranging from 375 m to 1 km per pixel [12,13]. This is substantially lower than the centimeter-level precision offered by UAV-driven data. Additionally, satellite datasets may contain data gaps or missing values (NaN) for certain attributes in specific regions due to cloud cover or sensor limitations. Consequently, relying exclusively on satellite data for detecting fires, particularly during the incipient stages, remains a challenging task.

Having weighed the respective trade-offs of UAV- and satellite-based systems, this study seeks to address two critical questions: How can the distinct strengths of these platforms be integrated to establish a more robust and comprehensive framework—one that not only detects wildfires in their initial stages but also accurately forecasts their daily progression dynamically? Furthermore, can UAVs serve as a complementary data source to enhance the precision of satellite-based perimeter predictions?

Building upon our previous work [14], this research presents a proof-of-concept UAV–satellite hybrid pipeline for power-efficient wildfire detection and dynamic perimeter prediction that synergizes the distinct advantages of UAV and satellite platforms. The pipeline operates in sequential phases, beginning with a multi-UAV search mission within an Area of Interest (AoI) identified as high-risk based on historical satellite data. Upon detection of a wildfire, the UAV transmits critical local data—including the detection timestamp, geocoordinates, and initial burned area—to a central server. This real-time input is then fed into an ML model, trained on a custom dataset developed for Canadian wildfire patterns, to generate a prediction of the fire's perimeter for the subsequent day. Conceptually, as the UAVs continue to relay updated on-site information, these predictions are dynamically refined. The primary contributions of this work are summarized below:

- Developed a hybrid UAV–satellite pipeline that utilizes a low-power, edge-computing model for onboard wildfire detection at the incipient stage, while leveraging satellite data to generate comprehensive perimeter predictions.
- Designed an interactive fire likelihood dashboard, derived from two decades of Canadian historical fire data (2000–2020), to provide a fire risk metric for any specific area across Canada.
- Curated a specialized dataset tailored to Canadian topography and vegetation, specifically structured to train models for fire perimeter and burned-area prediction.
- Developed a dynamic forecasting framework using a temporal simulation scheme. By sequentially feeding the model with daily fire status updates—mimicking real-time UAV data acquisition—we demonstrated the feasibility of iteratively refining perimeter predictions for the subsequent day as the fire evolves.

The remainder of this paper is organized as follows: Section 2 provides a comprehensive review of recent literature regarding wildfire detection and prediction methodologies. Section 3 details the architecture and operational phases of the proposed hybrid pipeline. Section 4 presents the curated dataset, pre-processing, and training details, along with an analysis of the chosen model architecture. Section 5 demonstrates experimental validation of the dynamic perimeter prediction framework, utilizing a simulation-based approach to establish a robust lower-bound performance baseline by evaluating model convergence against final cumulative burned masks in the absence of daily ground truth. Finally,

Section 6 concludes the manuscript by synthesizing the core findings of our proposed hybrid architecture.

2. Related Works

In recent years, the domain of wildfire detection and prediction has witnessed a significant proliferation of proposed methodologies, with a particular emphasis on the integration of machine learning techniques. The existing body of literature can be broadly categorized into two primary streams: UAV-based approaches and satellite-based systems.

In UAV-based methods, UAVs play the main role in the framework by providing the data for the ML model to gain the desired output. Several recent studies focused on UAV-driven segmentation models for describing fires. Feng et al. [15] proposed U3UNet for precise UAV fire segmentation and demonstrated high accuracy across complex environmental conditions. Similarly, Pesonen et al. [16] showed that compact neural networks could achieve near-real-time segmentation onboard drones, while Marinaccio and Afghah [17] demonstrated that RGB imagery alone could approximate temperature mapping through knowledge distillation. A temporal semantic segmentation approach was proposed by Meleti and Razi [18] that exploits changes in smoke patterns to detect obscured fire flames from UAV imagery. Jonnalagadda and Hashim [19] proposed a segmentation-based CNN to improve wildfire detection from UAV images that focused on reducing computational load and enhancing detection accuracy by segmenting high-resolution imagery into critical fire features. A substantial body of work emphasized fire detection and recognition. Zhu et al. [20] developed an enhanced YOLOv8 model for wildfire and smoke detection in complex forest environments from UAV data that improved feature extraction with attention modules to reduce false positives caused by vegetation and lighting changes. Li et al. [21] developed Fire-Net, a multi-stage lightweight CNN with cross-channel attention for UAV forest fire recognition that provides fast inference while maintaining high accuracy. A multi-UAV wildfire monitoring system is presented by Bailon-Ruiz et al. [22] that introduces trajectory planning for a fleet of fixed-wing UAVs to monitor wildfire progression over time. Some other UAV-based wildfire studies focused on advancing lightweight and accurate aerial perception models. One work proposed a modified Miti-DETR transformer–CNN hybrid model for UAV-based wildfire detection [23] and achieved precise localization with reduced computational overhead, enabling efficient onboard deployment. In [24], the authors proposed an aerial-image-based wildfire detection approach using color and multi-color space local binary pattern features of flame and smoke combined with a single neural network classifier. Their proposed model outperformed traditional classifiers and YOLOv3, demonstrating robustness for challenging conditions such as small flames, varying illumination, and complex backgrounds. Another study [25] employed an optimized YOLOv5m network with feature fusion and transfer learning to detect early wildfire smoke from UAV imagery and demonstrated improved sensitivity to small smoke targets. A separate effort developed lightweight fire and smoke detection methods using DeepLabv3+ and mobile vision transformers tailored for UAV constraints, showing that compact architectures could maintain strong segmentation performance while remaining efficient [26]. Finally, the DeepFire S3GA-Net framework [27] introduced a multiscale encoder–decoder segmentation architecture for UAV wildfire monitoring and achieved high segmentation accuracy across diverse environments while robustly extracting fire regions.

Recent satellite-based wildfire research has increasingly focused on improving active fire detection, burned-area segmentation, and near-real-time monitoring using multispectral and thermal space-borne sensors. Several studies focused on enhancing active fire detection by combining high-resolution polar-orbiting data with geostationary imagery, where approaches integrating VIIRS thermal anomaly products with Himawari-8 demonstrated

improved sensitivity to smaller ignition events and greater robustness to smoke and atmospheric interference [28,29]. Comparative work, such as Karlsson et al. [30], showed that VIIRS-based datasets provided superior predictability and reliability compared to MODIS, establishing the value of higher-spatial-resolution fire masks. Other studies evaluated large-scale Landsat-8 and Sentinel-2 satellite data, as in Pereira et al. [31] and Al-Dabbagh et al. [32], where deep learning approaches successfully identified and segmented active fires with meaningful improvements over classical spectral thresholding methods. Satellite-based burned-area mapping also received significant attention. Multiple works exploited Sentinel-2 reflectance data and spectral indices such as Normalized Burn Ratio (NBR) and Normalized Difference Vegetation Index (NDVI) to delineate burned regions, with studies such as Mahmood et al. [33], Tiengo et al. [34], and Liu et al. [35] showing that multi-index change detection approaches produced accurate regional burned-area products. More recent deep learning segmentation research, including Sui et al. [36], introduced bi-temporal architectures that leveraged pre-fire and post-fire imagery to enhance boundary precision and temporal coherence, outperforming conventional single-image detection methods. Beyond static detection and mapping, several contributions addressed temporal monitoring and prediction. Spatial–temporal deep learning frameworks used VIIRS time-series data to estimate daily burned-area evolution and capture wildfire growth dynamics more consistently than conventional static approaches that rely only on hotspots [37,38]. Shadrin et al. [39] introduced a large-scale wildfire spread prediction by developing a MA-Net neural network that is trained to predict the burned areas in a 5-day horizon. They reported that the prediction performance varied based on the prediction horizon, and their feature importance analysis identified land cover and wind data as the most influential factors in wildfire spread prediction. In [40], the authors introduced a comprehensive historical dataset focused on next-day wildfire spread, spanning nearly a decade of fire events in the United States. To demonstrate the dataset's utility, they implemented a neural network for spread prediction, benchmarking its performance against various other machine learning models. Subsequently, Ref. [41] utilized this same dataset to develop a U-Net architecture incorporating attention blocks for forecasting next-day fire progression. Other works demonstrated that combining multispectral and SAR observations helped maintain reliable burned-area mapping performance under cloud-obscured conditions [42], while large-scale global burned-area monitoring and cloud-platform processing studies showed that automated continental-scale wildfire tracking had become feasible [43]. Multi-source environmental integration research highlighted that coupling satellite observations with vegetation condition, fuel metrics, and climatic drivers strengthened predictive capability and supported higher-resolution fire-risk estimation [44].

3. Proposed Pipeline

As illustrated in Figure 1, our wildfire detection and dynamic perimeter prediction pipeline is structured into five distinct phases. The process initiates at Phase 0, where a high-risk AoI is identified using historical fire data. Based on the spatial extent of the AoI, the central server allocates a specific number of UAVs for a cooperative search mission, defining initial trajectories that can be dynamically updated during flight. In Phase 1, the UAVs execute the surveillance mission in 'capture-and-analyze' mode, flying over the target area and periodically capturing imagery. These images are analyzed locally using a custom onboard TinyML fire detection model [14], trained and evaluated on two separate datasets [45,46]. The model performs binary classification (fire vs. no-fire); instances classified as fire are immediately transmitted to the base station for validation.

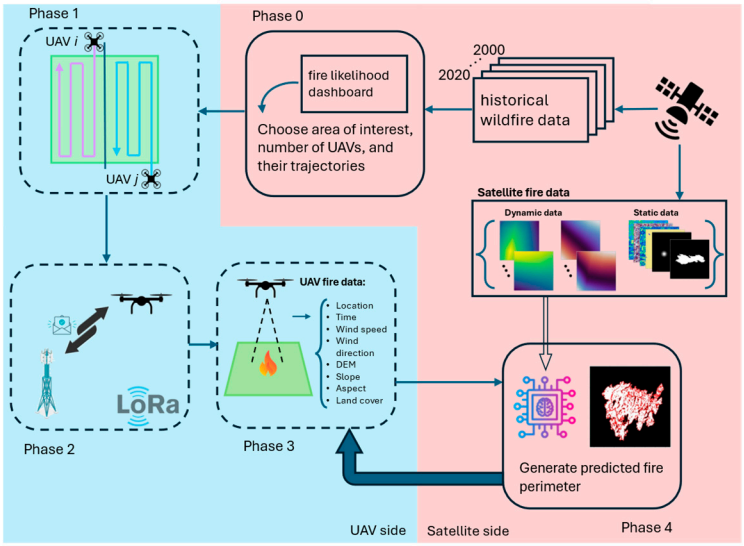


Figure 1. Proposed pipeline overview.

Given that the onboard TinyML model is optimized for power-constrained MCUs, its inference outputs require external validation to mitigate potential false positives. Therefore, in Phase 2, the candidate image transmitted by the UAV undergoes secondary analysis via a computationally robust ML model hosted on the central server. Upon verifying the presence of fire, the server transmits a confirmation signal back to the UAV. In Phase 3, triggered by this confirmation, the UAV transitions from capture-and-analyze mode to surveillance mode. In this state, it maneuvers over the active fire zone, capturing high-resolution imagery of flames, smoke, and the burned area, alongside real-time environmental telemetry. Following successful data transmission, the UAV returns to the base station for recharging or fleet rotation.

In the final phase, the data transmitted by the UAV is input into an ML model to generate a predicted burn map for the subsequent day. This perimeter prediction model is trained and evaluated on the specialized dataset detailed in Section 4. The UAV redeploys to the fire zone at server-determined intervals to provide updated status information, including imagery of the newly burned area. During these revisit missions, the server dynamically adjusts the UAV's trajectory to track the fire's propagation. The incoming data is then used to validate the previous perimeter prediction and update the model, yielding a refined forecast for the following day. This iterative process enables dynamic fire spread prediction, supporting timely and informed intervention strategies.

3.1. Phase 0: Initialization

While the proposed pipeline is deployable in any remote region, using Streamlit v1.55 we have developed a dedicated graphical dashboard to quantify fire likelihood for specific areas within Canada, leveraging historical Canadian wildfire data [47]. This likelihood metric (L_z) is formulated as a weighted aggregate of the zone's historical fire activity relative to the provincial totals. Specifically, it considers the frequency (N), magnitude (M, defined as burned area), and duration (T) of past fire events. The risk score is calculated as follows:

$$L_z = w_n\left(\frac{N_z}{N_P}\right) + w_m\left(\frac{M_z}{M_P}\right) + w_t\left(\frac{T_z}{T_P}\right), \quad (1)$$

where N_z, M_z, and T_z represent the total number of fires, total burned area, and cumulative duration within the selected zone, respectively. These values are normalized by the corresponding provincial totals (N_P, M_P, T_P). The coefficients w_n, w_m, and w_t are

tunable weighting factors satisfying $w_n + w_m + w_t = 1$, allowing the system to prioritize specific risk dimensions. Figure 2 illustrates a sample likelihood assessment for a region in British Columbia using this metric. Based on the spatial dimensions of the target area, the system determines the optimal number of UAVs and generates their corresponding search trajectories.

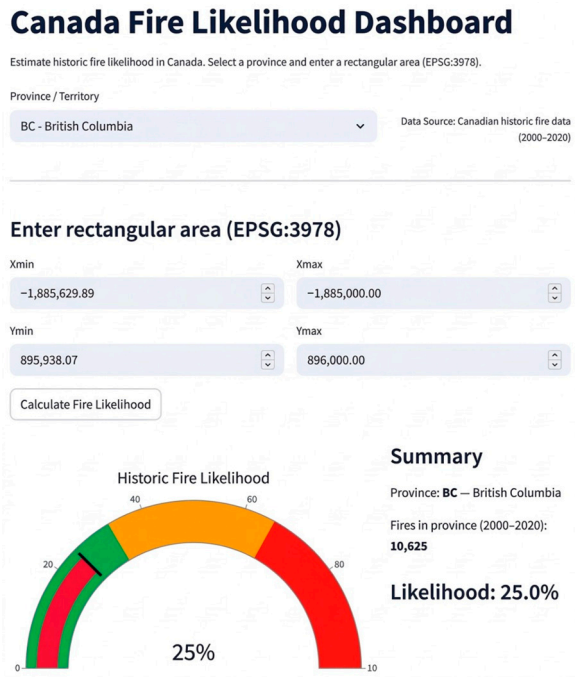


Figure 2. Fire likelihood dashboard: an example area in BC.

For the search mission, we assume a homogeneous fleet of UAVs operating at a constant altitude relative to ground level (AGL) and maintaining a uniform cruising speed. The AoI is partitioned equally among the agents, resulting in uniform power consumption rates and an identical maximum flight duration, denoted as τ_{max}. We further assume that the deployment latency—the time required for UAVs to reach their initial search coordinates—is negligible. The total operational time is discretized into n time steps. Consequently, the ground area coverage A_i captured by the i-th UAV in a single time step can be defined as a function of the flight altitude (h) and the camera's horizontal and vertical Field of View (FoV) angles (θ and ϕ). As illustrated in Figure 3, this coverage footprint is expressed as (2) and (3):

$$W_i = 2htan\left(\frac{\theta}{2}\right), \tag{2}$$

$$H_i = 2htan\left(\frac{\phi}{2}\right), \tag{3}$$

where i is the UAV counter and $i = \{1, \ldots, U\}$. Assuming symmetric FoV ($\theta = \phi$), have the following:

$$A_i = 4h^2tan^2\left(\frac{\theta}{2}\right), \tag{4}$$

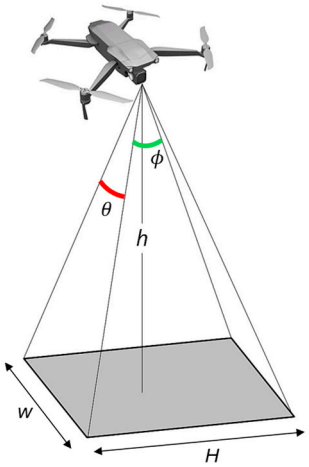


Figure 3. UAV's FoV angles and covered area.

However, to reflect real-world conditions, Equation (4) can be written as:

$$A_{i,eff} = \eta(1-\alpha)A_i = \eta(1-\alpha)4h^2 tan^2\left(\frac{\theta}{2}\right), \tag{5}$$

where η is the coverage efficiency, which is the ratio of actual covered area to ideal geometric area ($0 < \eta \leq 1$), and α is the coverage overlap.

The total operational flight duration is modeled as the cumulative sum of the time consumed across all n discrete steps of the search missions:

$$\tau_{max} = \sum_{k=1}^{n} t_k = \sum_{k=1}^{n} \left(t_{inf} + t_m + c.t_{turn}\right)_k, \tag{6}$$

where t_k represents the total duration of the k-th time step. This interval comprises three distinct components:

1. t_{inf}: The time required to capture imagery of the effective coverage area $A_{i,eff}$ and execute the onboard TinyML inference to detect fire presence.
2. t_m: The transit time needed to traverse the horizontal distance W along the trajectory to reach the center of the subsequent coverage zone.
3. $c.t_{turn}$: The maneuvering time associated with changing direction, where t_{turn} is the duration of a turn and $c \in \{0, 1\}$ is a binary indicator variable that equals 1 if a turning maneuver is executed at step k, and 0 otherwise.

The aggregate effective area covered by a single UAV over its maximum flight duration (τ_{max}) is given by $nA_{i,eff}$. Consequently, the minimum fleet size (U) required to completely survey the AoI is calculated as (7):

$$U = \left\lceil \frac{S}{nA_{i,eff}} \right\rceil = \left\lceil \frac{S}{n\eta(1-\alpha)4h^2 tan^2\left(\frac{\theta}{2}\right)} \right\rceil, \tag{7}$$

where S denotes the total surface area of the AoI. Following the determination of U, the AoI is decomposed into U congruent sub-regions, with each agent assigned to a specific partition. The UAV initializes its mission at the outer boundary of its designated zone and executes a systematic boustrophedon (lawn-mower) scanning pattern, as illustrated in Phase 1 of Figure 1. To optimize energy efficiency, the sweep direction is aligned parallel to the sub-area's longest edge; this geometric heuristic minimizes the number of energy-intensive turning maneuvers required to complete the coverage.

3.2. Phase 1: UAV Fire Detection

As established in our preceding study [14], effective wildfire detection can be achieved using a TinyML model deployed on a low-power, lightweight MCU. This approach significantly reduces the computational energy budget, thereby preserving battery life and extending operational flight duration. Adopting this strategy, our pipeline initiates the search mission (Phase 1) with UAVs executing their pre-determined trajectories, as illustrated in Figure 1. The agents operate in a continuous capture-and-analyze mode, where aerial imagery is sequentially acquired and processed locally by the onboard inference engine. Upon classifying an image as fire, the system triggers an immediate transmission of the candidate frame to the central server for further validation. The fire detection model is trained and evaluated utilizing two distinct datasets [45,46]. Table 2 presents a comparative analysis of the model's performance metrics relative to prominent state-of-the-art approaches.

Table 2. Comparative performance analysis of the adopted TinyML model against state-of-the-art fire detection methods.

Method	Accuracy	Precision	Recall	F1-Score	Dataset
U3UNet [15]	N/R *	0.911	0.881	0.896	Customized FLAME
[16]	0.967	0.912	0.73	0.772	Custom dataset
[20]	N/R	0.885	0.857	0.871	Custom dataset
MITI-DETR [23]	N/R	0.862	0.821	0.841	FLAME
Multi-Color ANN [24]	N/R	0.89	0.80	0.84	FLAME
MobileViT-Deeplabv3+ [26]	N/R	0.950	0.961	0.933	Customized FLAME
Tinier-YOLO [48]	N/R	0.798	0.79	0.81	FLAME
Proposed Tiny ML	0.918	0.90	0.937	0.91	FLAME
	0.974	0.969	0.991	0.974	UAVs-FFDB

* N/R: Not reported in the original publication.

3.3. Phase 2: Server Validation

To mitigate potential misclassifications from the onboard model, images flagged as fire are transmitted to the central server for secondary verification. This server-side validation stage is deliberately designed as a flexible, modular component within the pipeline. As the server operates without the computational and memory constraints of the UAVs, it can utilize robust, high-capacity deep learning architectures—such as a standard ResNet-50—to rigorously validate the candidate imagery. To ensure reproducibility and consistent feature extraction across the pipeline, this server-side validator is intended to be trained on the same data distributions (i.e., FLAME, UAVs-FFDB, and other relevant datasets). Given that the highly quantized edge model on the UAV already achieves strong baseline accuracy, deploying a deeper, unconstrained network on the server acts as a robust secondary filter, significantly reducing the likelihood of residual false positives. If the presence of fire is confirmed by this validator, a signal is transmitted to the UAV. In the absence of confirmation, the UAV dismisses the alert and resumes its standard capture-and-analyze operations.

3.4. Phase 3: Fire Data Collection by UAV

Upon receipt of the confirmation signal, the UAV transitions to surveillance mode, maneuvering to capture critical initial data, including the detection timestamp, geospatial coordinates, and the burned-area geometry. This telemetry is immediately transmitted to the central server to facilitate the subsequent prediction phase. Following data offload, the UAV returns to the base station for recharging or fleet rotation. To facilitate continuous monitoring and model validation, the UAV redeploys to the fire zone at server-determined intervals (e.g., daily) to provide updated status information, including imagery of the newly burned area. During these revisit missions, the server dynamically adjusts the UAV's trajectory to track the fire's propagation.

UAV–Server Communication Link

Complementing the computational efficiency of the onboard TinyML model, optimizing the UAV–Server communication link is essential, as wireless transmission constitutes a significant portion of the overall energy budget. While Long-Range (LoRa) technology offers advantageous low-power, long-distance connectivity, it is constrained by limited bandwidth. This bottleneck results in prolonged transmission times, necessitating extended flight durations to complete data uplinks, which ultimately increases the total energy expenditure. To mitigate these bandwidth constraints, this work adopts the Dual LoRa communication framework proposed in [49]. This architecture exploits the quasi-orthogonality of LoRa Spreading Factors (SFs) to integrate two parallel radio modules within a single node. Coordinated by a resource scheduling algorithm, this design enables simultaneous multi-channel transmission. Figure 4 illustrates the laboratory prototype of the Dual LoRa node. Experimental results reported in [49] indicate that this configuration achieves data rate enhancements of up to 73.5% compared to conventional single-radio implementations, contingent on transmission parameters. Consequently, this scheme is selected for the proposed pipeline as it effectively balances the critical requirements of power efficiency and low-latency communication. Table 3 presents the comparative transmission performance for a 24 kB payload under Line-of-Sight (LOS) conditions. The study evaluates the proposed Dual LoRa architecture (operating at 915 MHz with a 250 kHz bandwidth, utilizing simultaneous SF of 7 and 8) against a standard single LoRa link configured with identical spectral parameters and a fixed SF of 7. As quantified in the table, the Dual LoRa architecture significantly enhances transmission efficiency compared to the single-link baseline. Specifically, the dual radio configuration reduces the transmission latency for a 24 kB payload by 34.27% (dropping from 29.138 s to 19.153 s). Notably, this throughput gain is achieved without a corresponding penalty in power consumption; the Dual LoRa node actually maintains a marginally lower aggregate power draw compared to the single LoRa setup, thereby demonstrating superior energy efficiency.

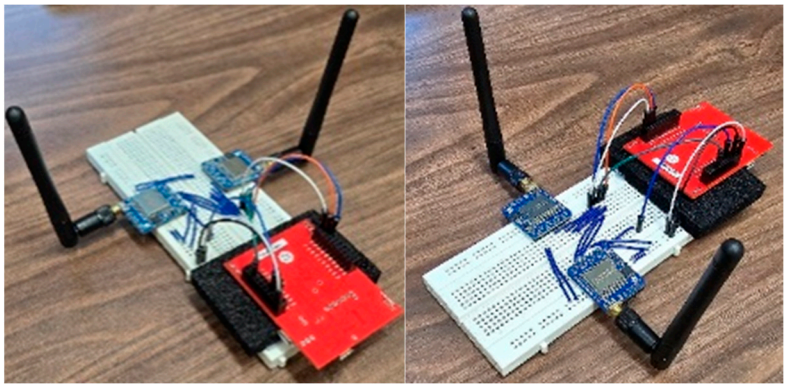


Figure 4. Lab prototype of a Dual LoRa node [49].

Table 3. Comparative results of data transmission over single LoRa versus Dual LoRa [49].

Wireless Link	SF	Bit Rate (kbps)	T_{tx} (s)	P_{tx} (mW)	P_{rx} (mW)	P_{total} (mW)
Single LoRa	7	822.71	29.138	287.25	37.75	325
Dual LoRa	7&8	1251.61	19.153	286.64	37.96	324.6

3.5. Phase 4: Fire Perimeter Prediction

In Phase 4, the data acquired by the UAV undergoes pre-processing to serve as the input for a machine learning model, which has been pre-trained on our satellite dataset. This model is tasked with forecasting the fire perimeter for the subsequent 24 h interval. To validate these predictions and close the feedback loop, the Phase 3 surveillance protocol is reiterated the following day to acquire ground-truth data of the actual burned area. By assimilating this fresh ground truth alongside updated environmental variables—specifically wind vectors and Leaf Area Index (LAI)—the model generates a refined perimeter forecast for the succeeding day. This iterative cycle establishes a hybrid, near-real-time framework for dynamic fire spread prediction.

4. Data Acquisition and Dataset Curation

This study specifically targets the unique dynamics of Canadian forest fires. While datasets have been established for other geographic regions, public benchmarks specifically tailored to fire perimeter and spread prediction within the Canadian forestry context appear to be limited in the current literature. To address this potential gap and ensure optimal data suitability for our application, we constructed a specialized dataset by synthesizing historical wildfire data (HWD) published by the Government of Canada [47] with publicly available multi-source satellite imagery.

The HWD repository aggregates comprehensive wildfire records across all Canadian provinces and territories, encapsulating key attributes such as temporal span (start and end dates), spatial magnitude (hectares), ignition coordinates, and final burned geometries (polygons). Each event is indexed by a unique identifier. Although the raw archive covers the period from 1917 to 2024, this study utilizes a subset from 2001 to 2020 to align with the temporal availability of essential satellite inputs, specifically land cover data. To further constrain model complexity, the dataset was refined to include only fire events with durations between 1 and 60 days, resulting in a total of 2656 instances. These selected fires exhibit a wide range of spatial scales, varying from less than 1 km^2 to over 150 km^2. The distributions for fire duration and size are illustrated in Figures 5 and 6, respectively.

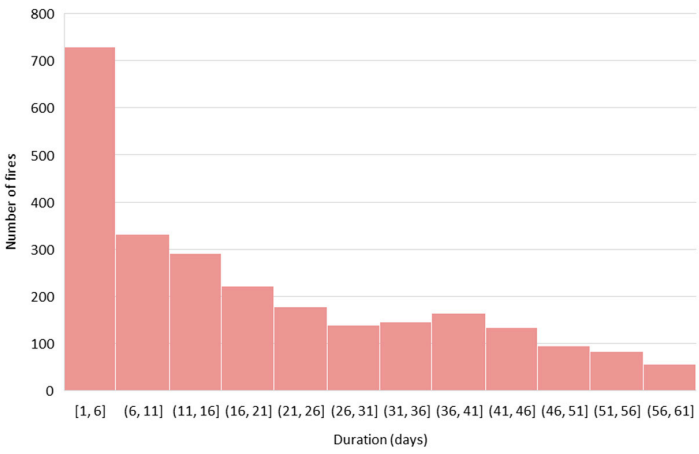


Figure 5. Distribution of fire durations.

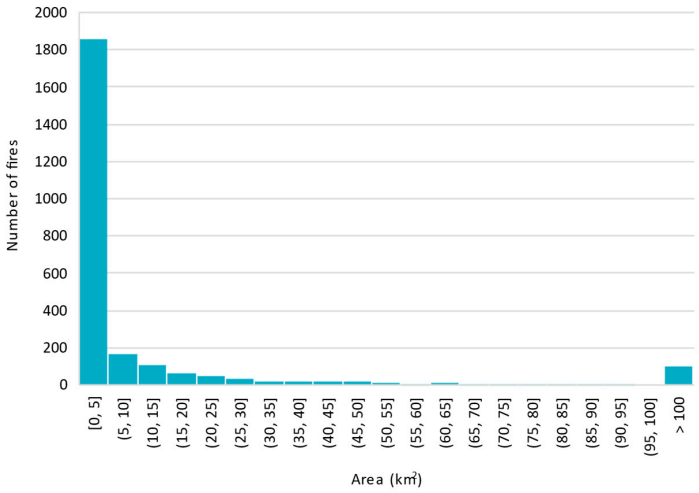


Figure 6. Distribution of burned-area sizes.

Table 4 provides a comprehensive specification of the public satellite resources utilized to construct the dataset. Topographic attributes—specifically elevation, slope, and aspect—are modeled as static features, premised on the assumption of geological stability over the two-decade study period (barring major geomorphic events such as landslides). Consequently, a single digital elevation model (DEM) acquisition is utilized for each Territory/Province (TP). In contrast, land cover is treated as a dynamic annual variable; thus, 20 distinct yearly layers were aggregated per TP to cover the 2001–2020 interval. To align with the model's objective of forecasting daily fire propagation, wind vector data (eastward and northward components) was selected at a daily temporal resolution. Finally, vegetation density is captured via Leaf Area Index (LAI) at an 8-day temporal resolution. Figure 7 illustrates a representative sample of these co-registered satellite inputs for a specific fire event from the HWD.

Table 4. Details of the multi-source remote sensing data utilized for dataset generation.

Feature	Description	Temporal Acquisition/Resolution	Spatial Resolution	Format	Data Source
Elevation (m)	Height of terrain	2011–2015/fixed	30 m	GeoTIFF	Copernicus GLO-30 digital elevation model [50]
Slope (degrees)	Steepness of terrain (0–90°)				
Aspect (degrees)	Direction of slope (0–360°)				
Land cover (classes)	Type of terrain	2001–2020/yearly	500 m		MODIS MCD12Q1 v061 [51]
Leaf area index (no unit)	Ratio of one-sided leaf area to unit ground area	2001–2020/8-day		NetCDF	MODIS MOD15A2H v061 [52]
Eastward wind (m/s)	Speed of eastward component of wind at 10 m above ground	2001–2020/daily	9 km		EAR5 [53]
Northward wind (m/s)	Speed of northward component of wind at 10 m above ground				

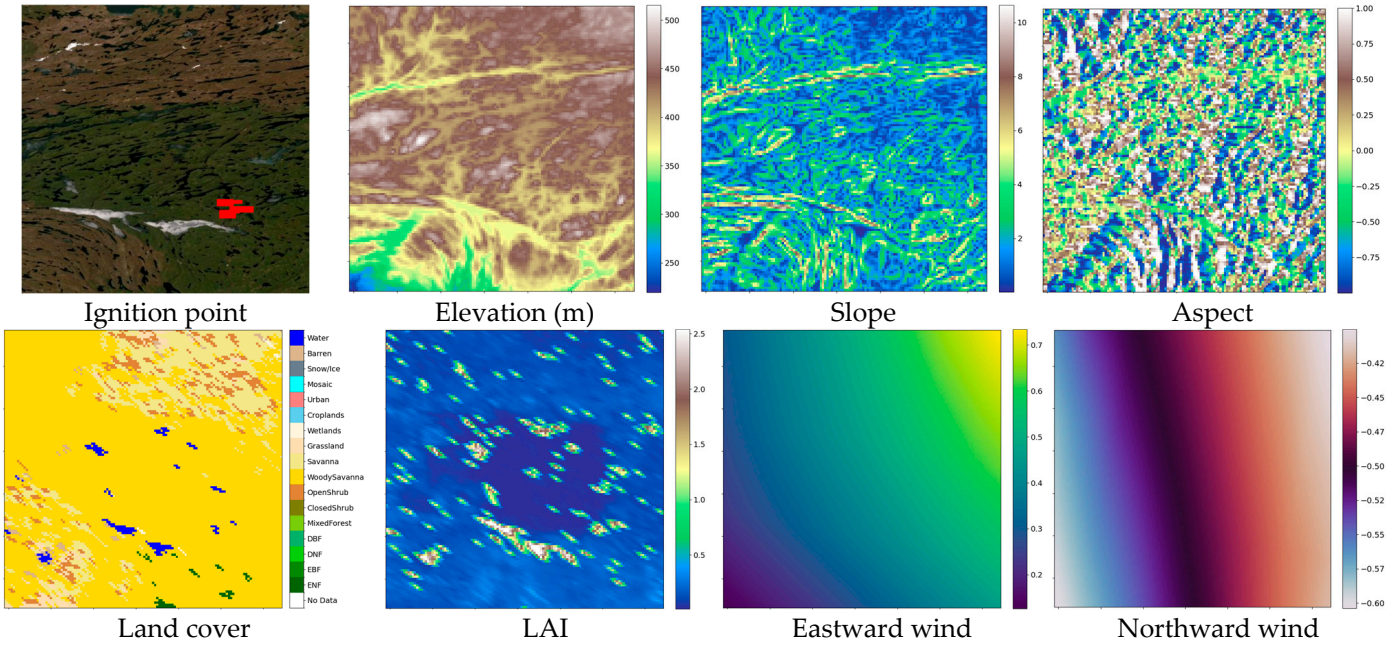


Figure 7. An example of satellite data for a random fire instance from HWD.

4.1. Data Pre-Processing and Training

For each individual fire, an adaptive spatial bounding box is generated, centered on the burned area's centroid. To capture sufficient immediate environmental context, the box's dimensions are initially scaled to exceed the fire's spatial extent by 30%, with a minimum baseline spatial window of 4 km established for smaller fires. During the pre-processing phase, the system evaluates the spatial grid of the ERA5 meteorological data. If the initial bounding box is too small to capture a meaningful meteorological gradient—defined as failing to encompass at least one ERA5 cell and its immediate neighbors—the box is systematically expanded until it successfully encapsulates the necessary wind data context. Once this optimized spatial extent is defined, the region of interest is cropped across all data sources. Finally, rather than forcing a rigid crop, the heterogeneous layers are smoothly interpolated using bilinear resampling into a uniform 128 $\times$ 128 pixel input tensor. The input volume is rigorously encoded to preserve feature integrity: the aspect feature is decomposed into sine and cosine components (two channels) to resolve angular cyclicity, while land cover is expanded into eight channels using one-hot encoding to represent its categorical nature without implying ordinality. Elevation and slope are retained as single channels. Defining fire duration (T) as the interval between the start and end dates, the LAI data is temporally aggregated into five distinct bins, resulting in five LAI frames per fire. For the eastward and northward wind components, which exist as daily sequences, consistency in input dimension is enforced by fixing the temporal depth to 60 days. Sequences shorter than this maximum duration are zero-padded for days beyond T, yielding 60 frames per wind component. Since the HWD is provided in the Canada Atlas Lambert (EPSG 3978) coordinate reference system, all satellite layers are reprojected to this standard prior to cropping. Furthermore, to convert the point-based ignition coordinates from the HWD into a spatial feature compatible with the ML model, a Gaussian blob is generated at the ignition locus with a variance proportional to the burned area. Figure 8 illustrates the resulting ignition map for the representative fire sample shown in Figure 7. Consequently, the total input depth per fire instance aggregates to 138 channels (13 static + 5 LAI + 120 wind).

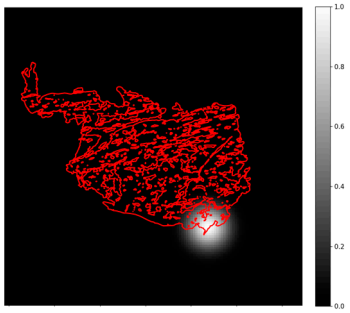


Figure 8. Gaussian blob of ignition point (white ball) drawn on the fire perimeter (red line).

Due to the strategic padding of bounding boxes to capture environmental context, the target masks exhibited a significant class imbalance, with background pixels heavily outnumbering the burned areas. To counteract this disparity and direct the model's focus toward critical growth patterns, the model was optimized using a spatially weighted Binary Cross-Entropy with Logits loss function (BCEWithLogitsLoss). This scheme calculates the standard pixel-wise BCE loss (reduction = 'none') and modulates it using a pre-computed weight map that assigns variable importance to four distinct topological zones:

1. Ignition anchor ($w = 50.0$): To ensure the model correctly identifies the fire's origin, the highest penalty is assigned to pixels where the ignition probability exceeds 0.4 within the ground-truth mask.
2. Fire core ($w = 15.0$): The interior of the active fire mask is weighted significantly higher than the background to prioritize recall.
3. Containment halo ($w = 10.0$): To penalize excessive false positives and force the model to respect the fire boundaries, a 'halo' ring is generated using morphological dilation. This zone is defined as a three-pixel-wide band situated two pixels outside the fire edge (created by subtracting a 2-iteration dilation from a 5-iteration dilation).
4. Background and buffer ($w = 1.0$): The remaining background pixels, along with a narrow 2-pixel 'buffer' zone immediately surrounding the fire, are assigned a baseline unit weight. This buffer allows for minor boundary uncertainties without harsh penalization.

The customized dataset was partitioned into a 70:15:15 split for training, validation, and testing, respectively. For the training of the perimeter prediction model, we utilized the AdamW optimizer with an initial learning rate of 10^{-4} and a weight decay of 10^{-3}. To facilitate optimal convergence, a ReduceLROnPlateau learning rate scheduler was applied. The model was trained with a batch size of four for a maximum of 100 epochs. To prevent overfitting, an early stopping mechanism with a patience of 10 epochs was implemented to monitor the validation loss.

Finally, to guarantee exact reproducibility of the results presented in the manuscript, a fixed random seed (42) was employed for data shuffling and weight initialization. Furthermore, we conducted multiple independent training runs utilizing different random seeds. These supplementary trials yielded consistent overall performance metrics and inference outputs, confirming that the model's predictive capability is robust, stable, and not an artifact of a specific weight initialization.

4.2. Model Architecture

To identify the optimal configuration for fire perimeter prediction, we conducted a comparative analysis of several prominent segmentation architectures, including U-Net, U-Net++, and DeepLabv3. These models were trained in combination with various loss functions—specifically Binary Cross-Entropy with Logits (BCEWithLogitsLoss), focal loss, and Tversky loss—to isolate the most effective objective function. To ensure fair

benchmarking, all experimental trials utilized the exact same training and validation splits. The resulting model performance was rigorously evaluated across a suite of standard metrics: F1-score, precision, recall, Intersection over Union (IoU), and Area Under the Curve (AUC).

Guided by the comparative performance analysis presented in Table 5, we selected the Squeeze-and-Excitation Residual U-Net (SE-ResUNet) as the optimal architecture for this application. This framework synergizes the residual learning strategy of ResNet [54] and the channel-attention mechanisms of Squeeze-and-Excitation networks [55] within a standard U-Net [56] encoder–decoder structure. The network design features an encoder with three downsampling stages leading to a bottleneck, where feature channels progressively expand from 64 to 128, 256, 512, and finally 1024. The decoder executes a symmetric reconstruction via skip connections, sequentially reducing channel depth from 512 back to 64. To enhance the model's sensitivity to relevant features, SE blocks are integrated into every residual unit, facilitating adaptive channel-wise recalibration. As evidenced in Table 5, the combination of this SE-ResUNet architecture with the BCEWithLogitsLoss function yielded superior segmentation metrics compared to all other evaluated configurations. Figure 9 presents the schematic of the SE-ResUNet architecture.

Table 5. Average F1-score across all prediction days and all test samples for different architectures and loss functions.

	UNet	UNet++	DeepLabv3	SE-ResUNet
BCE loss	0.59	0.59	0.61	0.63
Focal loss	0.60	0.62	0.62	0.62
Tversky	0.63	0.64	0.65	0.65
BCEwithLogits loss	0.62	0.64	0.64	0.66

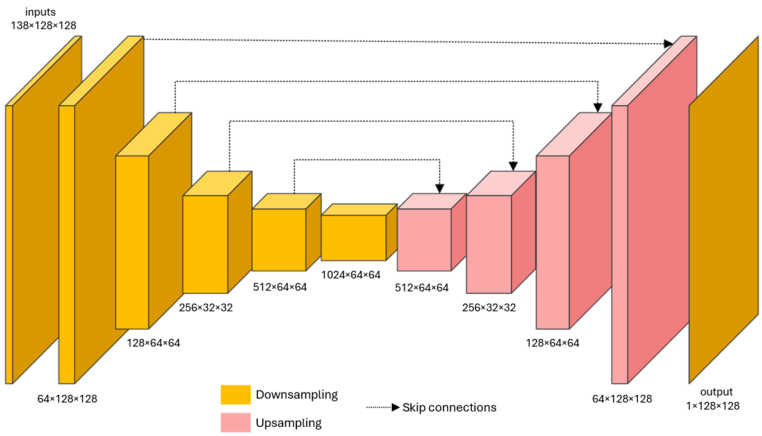


Figure 9. Schematic of SE-ResUNet model architecture used for fire perimeter prediction.

Table 6 presents the perimeter prediction performance metrics, calculated as an average across all test samples and daily prediction intervals, comparing the proposed pipeline's last component against related wildfire studies. The SE-ResUNet implementation achieves a recall of 0.84 $\pm$ 0.032, indicating a strong capability to capture the extent of the fire perimeter. In terms of overall segmentation accuracy, the module yields an F1-score of 0.66 $\pm$ 0.032, performing comparably to the highest benchmarks in the domain [39]. Additionally, the high AUC of 0.97 suggests the model effectively discriminates between burned and unburned classes across varied conditions. While the IoU of 0.51 reflects the challenge of exact boundary delineation in complex Canadian terrain, these results

effectively validate the proposed architecture, demonstrating the system's capacity to generate reliable perimeter forecasts as the final deliverable of the pipeline.

Table 6. Comparison of perimeter prediction performance metrics averaged over all test samples and daily prediction intervals.

Method	Precision	Recall	F1-Score	AUC	IoU
[30]	N/R *	N/R	0.272 ± 0.0063	N/R	0.16 ± 0.0042
[39]	N/R	N/R	0.66 ± 0.006	N/R	0.52
[40]	0.621	0.631	0.625	N/R	N/R
[41]	0.353	0.542	0.37	N/R	0.158
[44]	N/R	N/R	0.603	0.57	N/R
SE-ResUNet	0.65 ± 0.056	0.84 ± 0.032	0.66 ± 0.032	0.97 ± 0.001	0.51 ± 0.033

* N/R: Not reported in the original publication.

As detailed in Table 6, the perimeter prediction metrics exhibit a clear trade-off between recall (0.84) and precision (0.65). This tendency to over-predict burned areas is a deliberate architectural and operational outcome rather than an unmanaged artifact. From a technical standpoint, wildfire datasets suffer from severe class imbalance, with unburned background pixels vastly outnumbering the active fire perimeter. To prevent the model from defaulting to the majority background class, the spatially weighted loss function employed during training inherently penalized false negatives (missed fire pixels) far more aggressively than false positives, naturally biasing the network toward a higher recall. In wildfire management, a fail-safe overestimation is operationally preferable to a fatal underestimation. However, to prevent the deployment of an impractical model that flags excessively large regions as burning, we actively governed this conservative bias by shifting the inference decision threshold from the default 0.5 to a stricter 0.6. This specific calibration curbs runaway over-prediction by requiring the model to be at least 60% confident before classifying a pixel as fire, thereby securing a balanced and practical predictive output.

5. Dynamic Perimeter Prediction Simulation Results

To systematically evaluate the functional capability of the SE-ResUNet to forecast fire progression, we implemented a dynamic, simulation-based inference protocol. While the model is trained utilizing final cumulative burned maps due to the inherent scarcity of daily ground-truth maps in our dataset, this simulation is specifically designed to validate the network's spatiotemporal coherence. By demonstrating that the model's iterative daily predictions—driven by sequentially evolving environmental inputs—progressively converge with the final known burned geometry, we establish a robust indicator of its capacity to correctly interpret and propagate dynamic fire behavior over time.

Since in a real-world dynamic prediction environment, future temporal information is unavailable, and data is progressively acquired, we implemented an iterative inference protocol. Starting with the initial state (Day 0), we feed the model the static features alongside the first day's temporal data—while masking future temporal steps with zero-padding—to forecast the perimeter for Day 1. Subsequently, we incrementally reveal the next day's environmental data, accumulating the temporal history to predict the perimeter for Day 2. This sequential process is repeated until the final fire duration (T) is reached, effectively simulating the daily evolution of the fire perimeter. As illustrated in Figure 10, the architecture decouples the invariant static features (e.g., DEM, ignition) from the time-variant dynamic inputs. The control loop ($i = i + 1$) represents the sequential unmasking

of daily wind and LAI tensors, allowing the SE-ResUNet to generate a temporal sequence of segmentation maps that mirrors the fire's trajectory over T days.

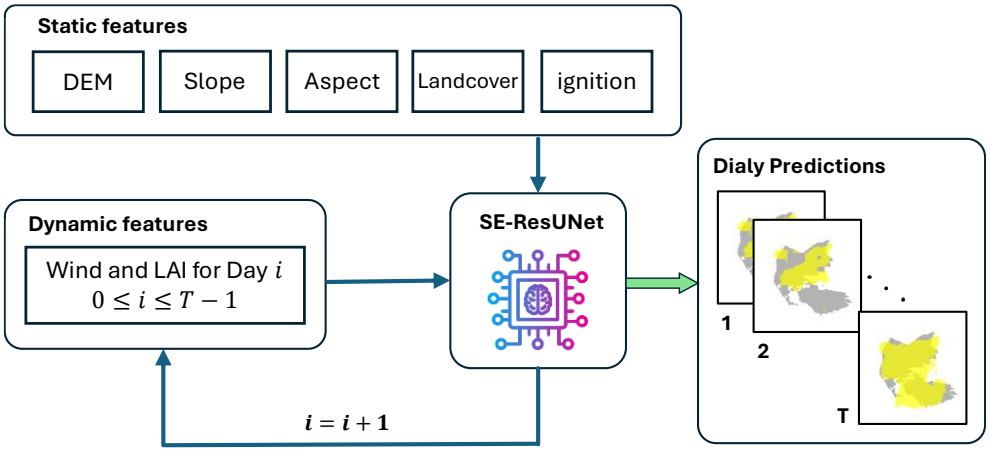


Figure 10. Schematic of the dynamic perimeter prediction simulation. The workflow demonstrates the iterative integration of daily environmental data (wind and LAI) with static terrain features to generate sequential fire progression maps.

The practical efficacy of the dynamic simulation framework is demonstrated in Figures 11 and 12, which visualize the daily perimeter evolution for two distinct fire events characterized by differing spatial scales and durations. It is important to contextualize the metric progression shown in these figures. A lower F1-score during the initial post-ignition phase does not necessarily indicate poor model performance or an inability to capture fire propagation. This behavior arises because daily ground-truth burned masks are not available, and model predictions at intermediate time steps are evaluated against the final burned-area mask. As a result, early predictions are penalized for not matching regions that will burn later in the event, introducing a bias in the evaluation of short-term forecasts. As additional temporal inputs become available, the model's predicted perimeter progressively aligns with the final burned mask. To measure the influence of daily burned masks on the performance of daily perimeter prediction, a separate experiment is conducted in Section 5.1 to simulate a live-fire scenario.

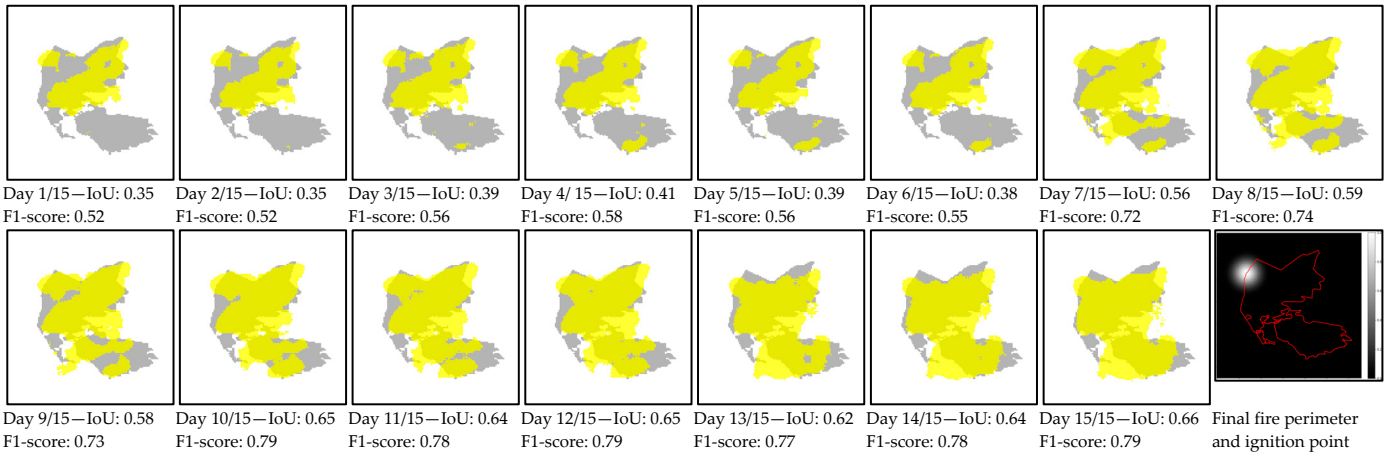


Figure 11. The performance of the proposed model on the historical dataset (15-day fire in Saskatchewan with a final burned area of 52 km^2). Yellow denotes the predicted perimeter; gray denotes the final burned map.

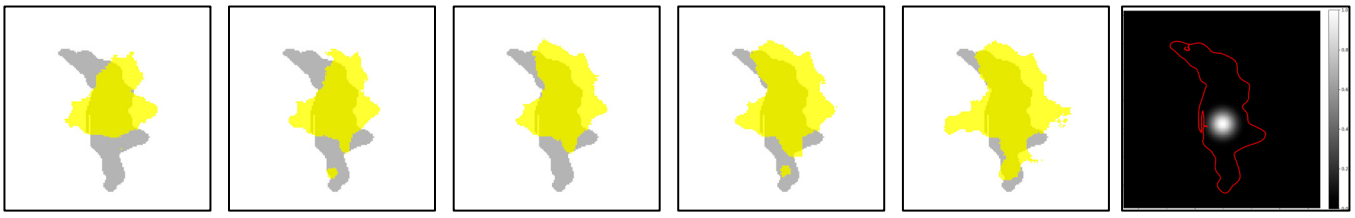


Figure 12. The performance of the proposed model on the historical dataset (5-day fire in Manitoba with a 7 km^2 size). Yellow is the predicted perimeter, and gray is the final burned map.

5.1. Impact of UAV Data on Daily Perimeter Prediction

An experiment is conducted to show the impact of having real-time ground data (new and updated imagery of the burned area) from the UAV on the proposed daily perimeter prediction model. Due to adverse seasonal conditions and strict regulatory constraints regarding prescribed burns, live field validation was not feasible within the scope of this research. Therefore, we conducted a simulation experiment using a subset of unseen fire instances from the test set. This experiment incorporates synthetically generated daily ground data as if they were collected from the UAV, which is then fed to the model to generate the predicted daily burned maps. This was achieved by applying a cost-distance propagation algorithm originating from the known ignition point. The propagation cost at each pixel was inversely weighted by a fusion of static topographic features (DEM, slope, and land cover) and dynamic environmental variables (daily wind vectors). By thresholding the cumulative cost surface, we produced a daily expansion sequence that adheres to the physical constraints of the local terrain and atmospheric conditions.

As illustrated in Figure 13, the Day 1 prediction is generated by providing the model with the primary UAV-derived ignition data alongside the satellite-based features and the current day's wind data. On Day 1, UAVs with updated trajectories are redeployed to the active fire location to capture images of the new burned area. This new data is replaced with the ignition to be input into the model along with Day 1 wind data to generate the Day 2 prediction. Overall, the simulated daily UAV-derived burned maps facilitate a closed-loop framework through three primary mechanisms. Correction: At the start of each simulated day, the model's internal state is reset by replacing its previous prediction with the latest UAV-derived map. Prediction: The model then uses this corrected state, combined with the current day's wind data, to predict the perimeter for the following day. Recursive update: This process is repeated daily, effectively allowing the model to produce predictions that are closer to reality, dynamically improving the daily forecast accuracy.

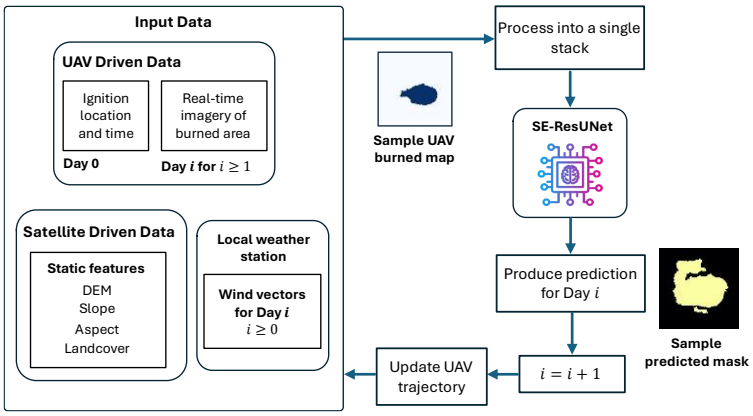


Figure 13. Flowchart of the simulation experiment for injecting UAV imagery of daily burned maps to predict fire on the next day in a live-fire scenario.

By injecting UAV spatial evidence into the model every day, the framework significantly reduces the drift between predicted and actual fire perimeters, as demonstrated in Figure 14. It demonstrates that, by incorporating real-time daily UAV observations, the framework effectively truncates error propagation, preventing the divergence in predictions that can occur by solely relying on historical satellite data. By re-initializing the model state with UAV-derived perimeters, the system maintains a high degree of spatial alignment, effectively correcting for potential drift in the predictive spread algorithm. Furthermore, it shows dynamic adaptation to localized terrain and fuel conditions while achieving a high precision–recall balance.

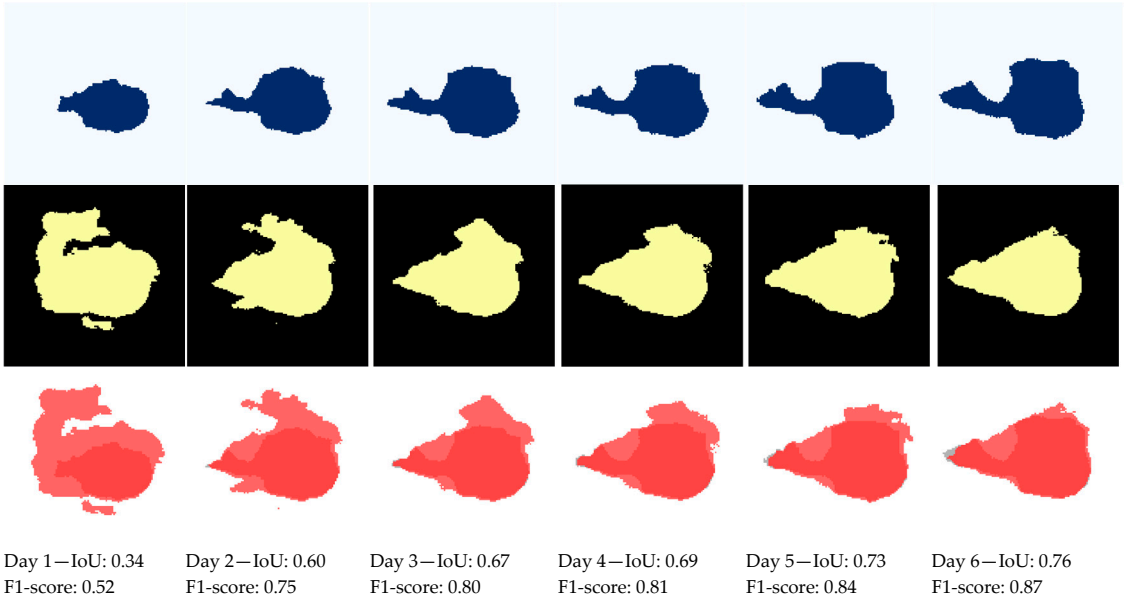


Figure 14. The performance of the prediction model on simulated real-time UAV maps and historical data. Top row: UAV real-time burned maps, middle row: predicted burned maps, and bottom row: overlaid burned maps.

Table 7 quantifies the impact of daily UAV-derived state updates on the average performance of the daily perimeter prediction framework across live-fire scenario samples (30% of the test set). It is evident that the integration of updated UAV perimeters via state injection leads to a significant increase in the mean IoU (+0.16) and F1-score (+0.13) compared to the same metrics for satellite-based predictions (Table 6). These results underscore that proximal UAV data is a critical component for achieving high-fidelity, real-time situational awareness. By providing high-frequency spatial evidence from the active fire front, the UAV enables the framework to maintain an accurate view of fire evolution, which is essential for reliable daily forecasting.

Table 7. Impact of simulated daily UAV burned maps on dynamic fire perimeter prediction performance metrics.

Data Source	Precision	Recall	F1-score	AUC	IoU
Satellite only	0.65 ± 0.056	0.84 ± 0.032	0.66 ± 0.032	0.97 ± 0.001	0.51 ± 0.033
Satellite + UAV	0.73 ± 0.034	0.91 ± 0.020	0.79 ± 0.049	0.98 ± 0.0005	0.67 ± 0.020

Table 8 provides a comprehensive comparative analysis, benchmarking the proposed pipeline against prominent state-of-the-art studies in the field. While existing works typically focus on isolated aspects of wildfire management—such as UAV-based detection

or satellite-based spread prediction—our framework uniquely integrates both domains into a cohesive system.

Table 8. Systematic comparison of the proposed pipeline against the existing wildfire literature, contrasting methodology, operational scalability, efficiency, and performance benchmarks.

Research	Application/ Method	Scalability	Power Efficiency	Dataset	Prediction Frequency	Overall Performance
[15]	Wildfire detection using UAV only	Single UAV deployment, Expandable dataset	-	Customized FLAME	-	Detection: F1-score: 89.6% Precision: 91.1% Recall: 88.1%
[23]	Wildfire detection using UAV only	Multiple UAV scalable	Low onboard computation overhead	FLAME	-	Detection: F1-score: 84.1% Precision: 82.1% Recall: 86.2%
[25]	Wildfire detection using UAV only	Single UAV deployment, Expandable dataset	Low onboard computation overhead	Customized FLAME	-	Detection: F1-score: 93.3% Precision: 96.1% Recall: 95%
[38]	Active wildfire detection and spread prediction using satellite only	-	-	Detection: FLAME and FLAME2 Prediction: custom dataset	Next day prediction	Detection: F1-score: 81.4% IoU: 70.9% Prediction: F1-score: 37.1% IoU 33.2%
[39]	Wildfire spread prediction using satellite only	-	-	Custom dataset	Daily up to 5 days	Prediction: F1-score: 66% IoU: 52%
[40]	Wildfire spread prediction using satellite only	Expandable dataset	-	Custom dataset	Next day prediction	Prediction: F1-score: 62.5%, Precision: 62.1%, Recall: 63.1%
[31]	Active wildfire detection using satellite only	-	-	Custom dataset	-	Detection: F1-score: 94.2% Precision: 92.9% Recall: 95.5%
[32]	Wildfire detection using satellite only	Expandable dataset	-	Custom dataset	-	Detection: F1-score: 98.24% Precision: 99.23% Recall: 97.32%
Proposed pipeline	Detection and dynamic perimeter prediction using UAVs and satellite data	Multiple UAV deployment, Expandable dataset	Low onboard computation overhead, Low-power communications	Detection: FLAME and UAVs-FFDB, Prediction: custom dataset	Prediction updated daily	Detection: F1-score: 91%, Precision: 90%, Recall: 93.7% Prediction: F1-score 66%, Precision: 65%, Recall: 84%, IoU: 51%

5.2. Discussion on Dataset and Limitations

The primary motivation for developing a customized dataset was the scarcity of publicly available benchmarks suitably tailored to the specific fire perimeter prediction challenges addressed in this research. To the best of our knowledge, existing repositories lacked the requisite spatiotemporal resolution or specific feature alignment necessary for the proposed Canadian forestry context. Following a rigorous review of the related literature, we identified and integrated the most critical environmental features governing perimeter evolution.

The designation of a UAV–satellite hybrid pipeline reflects the operational necessity of combining both platforms to overcome their respective dataset limitations. While satellite systems theoretically provide ignition locations, they suffer from significant temporal

latency and coarse spatial resolution, often failing to detect fires until they have already grown substantially. Therefore, the multi-UAV detection module is designed to provide near-real-time, high-confidence fire detection, serving as the critical, low-latency trigger that initializes the prediction phase.

Furthermore, the ultimate architectural design of our pipeline relies on the UAV to actively provide high-resolution, daily burned-area masks to iteratively update the prediction model as the fire evolves. Currently, the dataset's most notable limitation is the absence of these daily masks, which restricts the training capacity for fine-grained modeling of temporal fire dynamics. Capturing this specific experimental UAV data was unfeasible for this study due to strict aviation regulations and prohibitive winter weather conditions in Canada, which precluded prescribed fire operations. To bridge this operational gap, we conducted a simulation experiment demonstrating the impact of daily UAV-derived burned maps on prediction performance.

Beyond the absence of daily masks, the dataset retains other structural limitations. Incomplete temporal metadata—specifically missing end dates for certain events—results in indeterminate fire durations, obstructing the consistent extraction of temporal progression features. Additionally, data quality issues regarding missing or erroneous ignition coordinates affected a minority of samples, necessitating their exclusion from the training set. Furthermore, the integration of heterogeneous data sources characterized by disparate spatial resolutions introduces potential interpolation challenges. A primary example of this is the meteorological data; while the 9 km spatial resolution of the ERA5 dataset captures essential synoptic and mesoscale wind patterns, it cannot resolve micro-scale, topography-driven local wind channeling. Although our interpolation strategy ensures the data is mathematically smoothed and structurally aligned for the neural network, the model is inherently constrained by the physical absence of micro-local turbulence data. Nevertheless, for predicting daily (24 h) fire perimeter growth in the expansive Canadian boreal forest, mesoscale weather patterns remain the dominant climatic driver of fire progression.

Moreover, the strategic expansion of spatial context windows—necessary to capture these environmental surroundings—inherently exacerbates class imbalance. While this disparity is methodologically mitigated via the spatial weighting scheme, the underlying background dominance remains a structural characteristic of the data.

Finally, the dataset is intrinsically calibrated to the specific land cover, climate, and forest characteristics of the Canadian context. Consequently, the direct generalization of models trained on this data to regions with distinct fuel types, topography, or fire regimes may be constrained. Nevertheless, the proposed dataset establishes a robust foundation for perimeter prediction, offering a baseline that can be effectively extended to broader geographic contexts through regional domain adaptation and transfer learning strategies.

6. Conclusions

This paper presented a comprehensive UAV–satellite hybrid pipeline designed to address the dual challenges of real-time wildfire detection and dynamic perimeter prediction. By integrating the agility of aerial surveillance with the scalability of satellite-driven deep learning, the proposed framework establishes a modular and foundational architecture to support advanced disaster management. First, to identify high-risk areas of interest, we developed a graphical dashboard based on historical Canadian wildfire records that produces a fire-likelihood score for any selected region within Canada. Given the size of a selected area, the pipeline is initialized by allocating an appropriate number of UAVs and computing their trajectories. The UAVs begin their mission in a capture-and-analyze mode, where a TinyML model is executed onboard to detect early signs of fire. When an image is classified as fire, the result is validated at the server. Upon confirmation, the UAV

switches to surveillance mode, continuously transmitting localized observations of the detected event.

To ensure reliable connectivity while minimizing energy overhead, we implement a power-efficient Dual LoRa communication scheme. Experimental results show that the detection module achieves an F1-score of 91% on the FLAME dataset and 97.4% on the UAVs-FFDB dataset, enabling persistent, low-latency monitoring and mitigating the operational bottlenecks typically associated with continuous drone-based surveillance. To support the pipeline's predictive phase, we addressed the scarcity of region-specific benchmarks by creating a customized wildfire dataset tailored to the Canadian boreal context. This repository integrates static terrain features with dynamic environmental variables, providing a robust foundation for training deep segmentation models. Utilizing this data, we implemented a Squeeze-and-Excitation Residual U-Net (SE-ResUNet), which demonstrated superior performance in delineating fire fronts, achieving a recall of 84% and an AUC of 0.97.

Finally, we conducted an experiment to show the impact of having real-time burned-area imagery from the UAV on the proposed daily perimeter prediction model. Simulation results demonstrate that incorporating updated daily UAV observations effectively truncates error propagation, producing more accurate predictions compared to relying solely on historical satellite data.

Author Contributions: Conceptualization, H.K. and K.A.W.; methodology, H.K. and K.A.W.; software, H.K.; validation, H.K. and K.A.W.; formal analysis, H.K.; investigation, H.K.; resources, H.K.; data curation, H.K.; writing—original draft preparation, H.K.; writing—review and editing, H.K. and K.A.W.; visualization, H.K.; supervision, K.A.W.; project administration, K.A.W.; funding acquisition, K.A.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC), grant number RGPIN-2018-06274.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: During the preparation of this manuscript, the authors used Google Gemini 3 Flash, a 2026 LLM model, for the purposes of proofreading and text refinement. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kirchmeier-Young, M.C.; Malinina, E.; Barber, Q.E.; Perdomo, K.G.; Curasi, S.R.; Liang, Y.; Jain, P.; Gillett, N.P.; Parisien, M.-A.; Cannon, A.J.; et al. Human driven climate change increased the likelihood of the 2023 record area burned in Canada. *NPJ Clim. Atmos. Sci.* **2024**, *7*, 316. [CrossRef]
2. Kelley, D.I.; Burton, C.; Di Giuseppe, F.; Jones, M.W.; Barbosa, M.L.F.; Brambleby, E.; McNorton, J.R.; Liu, Z.; Bradley, A.S.I.; Blackford, K.; et al. State of Wildfires 2024–2025. *Earth Syst. Sci. Data* **2025**, *17*, 5377–5488. [CrossRef]
3. Thangavel, K.; Spiller, D.; Sabatini, R.; Amici, S.; Sasidharan, S.T.; Fayek, H.; Marzocca, P. Autonomous Satellite Wildfire Detection Using Hyperspectral Imagery and Neural Networks: A Case Study on Australian Wildfire. *Remote Sens.* **2023**, *15*, 720. [CrossRef]
4. Tong, H.; Yuan, J.; Zhang, J.; Wang, H.; Li, T. Real-Time Wildfire Monitoring Using Low-Altitude Remote Sensing Imagery. *Remote Sens.* **2024**, *16*, 2827. [CrossRef]
5. Fan, X.; Lei, F.; Yang, K. Real-Time Detection of Smoke and Fire in the Wild Using Unmanned Aerial Vehicle Remote Sensing Imagery. *Forests* **2025**, *16*, 201. [CrossRef]
6. Seydi, S.T.; Saeidi, V.; Kalantar, B.; Ueda, N.; Halin, A.A. Fire-Net: A Deep Learning Framework for Active Forest Fire Detection. *J. Sensors* **2022**, *2022*, 8044390. [CrossRef]
7. Singh, H.; Ang, L.-M.; Srivastava, S.K. Active wildfire detection via satellite imagery and machine learning: An empirical investigation of Australian wildfires. *Nat. Hazards* **2025**, *121*, 9777–9800. [CrossRef]

8. Drusch, M.; Del Bello, U.; Carlier, S.; Colin, O.; Fernandez, V.; Gascon, F.; Hoersch, B.; Isola, C.; Laberinti, P.; Martimort, P.; et al. Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services. *Remote Sens. Environ.* **2012**, *120*, 25–36. [CrossRef]
9. Roy, D.P.; Wulder, M.A.; Loveland, T.R.; Woodcock, C.E.; Allen, R.G.; Anderson, M.C.; Helder, D.; Irons, J.R.; Johnson, D.M.; Kennedy, R.; et al. Landsat-8: Science and product vision for terrestrial global change research. *Remote Sens. Environ.* **2014**, *145*, 154–172. [CrossRef]
10. NASA Earthdata. Available online: https://www.earthdata.nasa.gov/data/projects/lance (accessed on 9 December 2025).
11. Natural Resources Canada. Available online: https://cwfis.cfs.nrcan.gc.ca/ha/nfdb (accessed on 9 December 2025).
12. Schroeder, W.; Oliva, P.; Giglio, L.; Csiszar, I.A. The New VIIRS 375m active fire detection data product: Algorithm description and initial assessment. *Remote Sens. Environ.* **2014**, *143*, 85–96. [CrossRef]
13. Giglio, L.; Schroeder, W.; Justice, C.O. The collection 6 MODIS active fire detection algorithm and fire products. *Remote Sens. Environ.* **2016**, *178*, 31–41. [CrossRef] [PubMed]
14. Farahdel, A.; Khan, A.H.; Keshmiri, H.; Wahid, K.A. TinySurveillance: An Extra Low-Power Event-Based Surveillance Method for UAVs. *IEEE Access* **2025**, *13*, 98542–98556. [CrossRef]
15. Feng, H.; Qiu, J.; Wen, L.; Zhang, J.; Yang, J.; Lyu, Z.; Liu, T.; Fang, K. U3UNet: An accurate and reliable segmentation model for forest fire monitoring based on UAV vision. *Neural Netw.* **2025**, *185*, 107207. [CrossRef]
16. Pesonen, J.; Hakala, T.; Karjalainen, V.; Koivumäki, N.; Markelin, L.; Raita-Hakola, A.-M.; Suomalainen, J.; Pölönen, I.; Honkavaara, E. Detecting Wildfires on UAVs with Real-time Segmentation Trained by Larger Teacher Models. *arXiv* **2024**, arXiv:2408.10843. [CrossRef]
17. Marinaccio, M.; Afghah, F. Seeing Heat with Color—RGB-Only Wildfire Temperature Inference from SAM-Guided Multimodal Distillation using Radiometric Ground Truth. *arXiv* **2025**, arXiv:2505.01638. [CrossRef]
18. Meleti, U.; Razi, A. Obscured Wildfire Flame Detection by Temporal Analysis of Smoke Patterns Captured by Unmanned Aerial Systems. *arXiv* **2023**, arXiv:2307.00104. [CrossRef]
19. Jonnalagadda, A.V.; Hashim, H.A. SegNet: A segmented deep learning based Convolutional Neural Network approach for drones wildfire detection. *Remote Sens. Appl.* **2024**, *34*, 101181. [CrossRef]
20. Zhu, W.; Niu, S.; Yue, J.; Zhou, Y. Multiscale wildfire and smoke detection in complex drone forest environments based on YOLOv8. *Sci. Rep.* **2025**, *15*, 2399. [CrossRef]
21. Li, S.; Han, J.; Chen, F.; Min, R.; Yi, S.; Yang, Z. Fire-Net: Rapid Recognition of Forest Fires in UAV Remote Sensing Imagery Using Embedded Devices. *Remote Sens.* **2024**, *16*, 2846. [CrossRef]
22. Bailon-Ruiz, R.; Bit-Monnot, A.; Lacroix, S. Real-time wildfire monitoring with a fleet of UAVs. *Rob. Auton. Syst.* **2022**, *152*, 104071. [CrossRef]
23. Muksimova, S.; Umirzakova, S.; Mardieva, S.; Abdullaev, M.; Cho, Y.I. Revolutionizing Wildfire Detection Through UAV-Driven Fire Monitoring with a Transformer-Based Approach. *Fire* **2024**, *7*, 443. [CrossRef]
24. Hossain, F.A.; Zhang, Y.M.; Tonima, M.A. Forest fire flame and smoke detection from UAV-captured images using fire-specific color features and multi-color space local binary pattern. *J. Unmanned Veh. Syst.* **2020**, *8*, 285–309. [CrossRef]
25. Mukhiddinov, M.; Abdusalomov, A.B.; Cho, J. A Wildfire Smoke Detection System Using Unmanned Aerial Vehicle Images Based on the Optimized YOLOv5. *Sensors* **2022**, *22*, 9384. [CrossRef]
26. Kumar, A.; Perrusquía, A.; Al-Rubaye, S.; Guo, W. Wildfire and smoke early detection for drone applications: A light-weight deep learning approach. *Eng. Appl. Artif. Intell.* **2024**, *136*, 108977. [CrossRef]
27. Akhyar, A.; Lee, J.; Shahar, N.; Saputro, A.H.; Zulkifley, M.A. Automated forest fire detection in ecological monitoring using enhanced deep learning networks. *Sci. Rep.* **2025**, *16*, 2039. [CrossRef] [PubMed]
28. Zhang, D.; Huang, C.; Gu, J.; Hou, J.; Zhang, Y.; Han, W.; Dou, P.; Feng, Y. Real-Time Wildfire Detection Algorithm Based on VIIRS Fire Product and Himawari-8 Data. *Remote Sens.* **2023**, *15*, 1541. [CrossRef]
29. Kang, Y.; Jang, E.; Im, J.; Kwon, C. A deep learning model using geostationary satellite data for forest fire detection with reduced detection latency. *GIScience Remote Sens.* **2022**, *59*, 2019–2035. [CrossRef]
30. Karlsson, J.; Xu, Y.; Berg, A.; Haglund, L. Comparing Next-Day Wildfire Predictability of MODIS and VIIRS Satellite Data. *arXiv* **2025**, arXiv:2503.08580. [CrossRef]
31. Pereira, G.H.D.A.; Fusioka, A.M.; Nassu, B.T.; Minetto, R. Active Fire Detection in Landsat-8 Imagery: A Large-Scale Dataset and a Deep-Learning Study. *arXiv* **2021**, arXiv:2101.03409. [CrossRef]
32. Al-Dabbagh, A.M.; Ilyas, M. Uni-temporal Sentinel-2 imagery for wildfire detection using deep learning semantic segmentation models. *Geomat. Nat. Hazards Risk* **2023**, *14*, 2196370. [CrossRef]
33. Mahmood, M.; Jumaah, H. NBR index-based fire detection using sentinel-2 images and GIS: A case study in mosul park, Iraq. *Int. J. Geoinformatics* **2023**, *19*, 67–74. [CrossRef]

34. Tiengo, R.; Merino-De-Miguel, S.; Uchôa, J.; Guiomar, N.; Gil, A. Burned Areas Mapping Using Sentinel-2 Data and a Rao's Q Index-Based Change Detection Approach: A Case Study in Three Mediterranean Islands' Wildfires (2019–2022). *Remote Sens.* **2025**, *17*, 830. [CrossRef]
35. Liu, P.; Liu, Y.; Guo, X.; Zhao, W.; Wu, H.; Xu, W. Burned area detection and mapping using time series Sentinel-2 multispectral images. *Remote Sens. Environ.* **2023**, *296*, 113753. [CrossRef]
36. Sui, T.; Huang, Q.; Wu, M.; Wu, M.; Zhang, Z. BiAU-Net: Wildfire burnt area mapping using bi-temporal Sentinel-2 imagery and U-Net with attention mechanism. *Int. J. Appl. Earth Obs. Geoinf.* **2024**, *132*, 104034. [CrossRef]
37. Zhao, Y.; Ban, Y. Near real-time wildfire progression mapping with VIIRS time-series and autoregressive SwinUNETR. *Int. J. Appl. Earth Obs. Geoinf.* **2025**, *136*, 104358. [CrossRef]
38. Zhao, Y.; Gerard, S.; Ban, Y. TS-SatFire: A Multi-Task Satellite Image Time-Series Dataset for Wildfire Detection and Prediction. *arXiv* **2024**, arXiv:2412.11555. [CrossRef] [PubMed]
39. Shadrin, D.; Illarionova, S.; Gubanov, F.; Evteeva, K.; Mironenko, M.; Levchunets, I.; Belousov, R.; Burnaev, E. Wildfire spreading prediction using multimodal data and deep neural network approach. *Sci. Rep.* **2024**, *14*, 2606. [CrossRef]
40. Huot, F.; Hu, R.L.; Goyal, N.; Sankar, T.; Ihme, M.; Chen, Y.-F. Next Day Wildfire Spread: A Machine Learning Dataset to Predict Wildfire Spreading from Remote-Sensing Data. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 1–13. [CrossRef]
41. Fitzgerald, J.; Seefried, E.; E Yost, J.; Pallickara, S.; Blanchard, N. Paying Attention to Wildfire: Using U-Net with Attention Blocks on Multimodal Data for Next Day Prediction. In *ICMI '23: Proceedings of the 25th International Conference on Multimodal Interaction*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 470–480. [CrossRef]
42. Vittucci, C.; Cordari, F.; Guerriero, L.; Di Sanzo, P. Design and evaluation of a cloud-oriented procedure based on SAR and Multispectral data to detect burnt areas. *Earth Sci. Inform.* **2025**, *18*, 322. [CrossRef]
43. Nelson, D.M.; He, Y.; Moore, G. Trends and applications in wildfire burned area mapping: Remote sensing data, cloud geoprocessing platforms, and emerging algorithms. *Geomatica* **2024**, *76*, 100008. [CrossRef]
44. Porta, H.; Dalsasso, E.; McCarty, J.L.; Tuia, D. CanadaFireSat: Toward high-resolution wildfire forecasting with multiple modalities. *arXiv* **2025**, arXiv:2506.08690. [CrossRef]
45. Mowla, N.; Asadi, D.; Tekeoglu, K.N.; Masum, S.; Rabie, K. UAVs-FFDB: A high-resolution dataset for advancing forest fire detection and monitoring using unmanned aerial vehicles (UAVs). *Data Brief* **2024**, *55*, 110706. [CrossRef]
46. Shamsoshoara, A.; Afghah, F.; Razi, A.; Zheng, L.; Fulé, P.; Blasch, E. Aerial imagery pile burn detection using deep learning: The FLAME dataset. *Computer Networks* **2021**, *193*, 108001. [CrossRef]
47. National Fire Database Fire Polygon Data. Available online: https://cwfis.cfs.nrcan.gc.ca/downloads/nfdb/fire_poly/current_version/NFDB_poly.zip (accessed on 31 October 2025).
48. Fang, W.; Wang, L.; Ren, P. Tinier-YOLO: A Real-Time Object Detection Method for Constrained Environments. *IEEE Access* **2020**, *8*, 1935–1944. [CrossRef]
49. Keshmiri, H.; Rahman, G.M.E.; Wahid, K.A. LoRa Resource Allocation Algorithm for Higher Data Rates. *Sensors* **2025**, *25*, 518. [CrossRef] [PubMed]
50. Copernicus GLO-30 Digital Elevation Model. Available online: https://portal.opentopography.org/raster?opentopoID=OTSDEM.032021.4326.3 (accessed on 28 October 2025).
51. MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V061. Available online: https://www.earthdata.nasa.gov/data/catalog/lpcloud-mcd12q1-061 (accessed on 28 October 2025).
52. MODIS/Terra Leaf Area Index/FPAR 8-Day L4 Global 500m SIN Grid V061. Available online: https://www.earthdata.nasa.gov/data/catalog/lpcloud-mod15a2h-061 (accessed on 28 October 2025).
53. ERA5-Land Post-Processed Daily Statistics from 1950 to Present. Available online: https://cds.climate.copernicus.eu/datasets/derived-era5-land-daily-statistics?tab=overview (accessed on 28 October 2025).
54. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
55. Hu, J.; Shen, L.; Sun, G. Squeeze-and-excitation networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 7132–7141.
56. Ronneberger, O.; Fischer, P.; Brox, T. U-net: Convolutional networks for biomedical image segmentation. In Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, Munich, Germany, 5–9 October 2015; pp. 234–241.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.