NASA/TM–20260007872

Proactive Wildfire Management: A Remote Sensing and Multimodal CNN-MLP Architecture for Ignition Risk Forecasting

**** All authors contributed equally to this memo. The names are listed alphabetically by last name.***

Tyler Chen
Saint Francis High School, Mountain View, CA

Kaitlyn Mendoza
Mater Brickell Academy, Miami, FL

Praneel Mukherjee
Rock Ridge High School & Academies of Loudoun, Ashburn, VA

Neel Navuduri
Oxford Academy, Cypress, CA

Arush Shangari
St. John's Preparatory School, North Reading, MA

Arya Sira
Cypress Woods High School, Cypress, TX

August 2026

NASA STI Program Report Series

The NASA STI Program collects, organizes, and provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published on both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA is a counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release of reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/
and select the "General" help request type.

NASA/TM–20260007872

Proactive Wildfire Management: A Remote Sensing and Multimodal CNN-MLP Architecture for Ignition Risk Forecasting

** All authors contributed equally to this memo. The names are listed alphabetically by last name.*

Tyler Chen
Saint Francis Highschool, Mountain View, CA

Kaitlyn Mendoza
Mater Brickell Academy, Miami, FL

Praneel Mukherjee
Rock Ridge High School & Academies of Loudoun, Ashburn, VA

Neel Navuduri
Oxford Academy, Cypress, CA

Arush Shangari
St. John's Preparatory School, North Reading, MA

Arya Sira
Cypress Wood Highschool, Cypress, TX

National Aeronautics and
Space Administration

Ames Research Center
Mountain View, CA

August 2026

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Acknowledgment

Team 4 would like to take this opportunity to express our sincere gratitude to everyone who contributed to the success of this project. First, we extend our deepest thanks to our team mentors, Dr. Krishna Kalyanam and Dr. Milad Memarzadeh, whose valuable guidance, insightful feedback, and constant support were instrumental in refining our work and steering this project to completion. We are also extremely grateful to NASA for providing this unique learning opportunity, which enabled us to build vital research skills and expand our technical knowledge. Finally, we wish to thank all our team members for their dedication, seamless collaboration, and willingness to go the extra mile throughout this effort.

Table of Contents

1 Abstract

As the frequency and intensity of wildfires increase, with fire seasons now starting earlier and ending later than they have over the past decades, current monitoring systems, such as lookout towers and satellites, are hindered by cloud cover, low-resolution imagery, and static data gaps that fail to track vegetation moisture levels fast enough to catch rapid pre-ignition changes. This report proposes a Machine Learning-enabled Wildfire Ignition Prediction framework that combines satellite monitoring with dynamic and high-resolution remote sensing from Unmanned Aerial Vehicle (UAV) swarms. The method would use multispectral and thermal data from the Landsat program to create a baseline for vegetation health, calculating a two-band Enhanced Vegetation Index (EVI2) and the moisture content of the vegetation. These inputs will later be fused with microscale UAV weather data, including thermal hotspots found through thick canopies, hyperspectral chemical signatures of pre-visual combustion, and local weather streams. The multispectral satellite, multispectral Light Detection and Ranging (LiDAR), and thermal data would then be processed through a Convolutional Neural Network (CNN), alongside a Multilayer Perceptron (MLP) for the micro-weather telemetry. The outputs of these networks would be fused into a single feature representation and passed through a final prediction network to generate real-time ignition risk scores and hotspot alerts. Model performance would be assessed using standard classification metrics, including a Receiver Operating Characteristic - Area Under the Curve (ROC AUC) and F1 score. This system would allow first responders to identify high-risk zones and intervene before ignition occurs, improving emergency response time compared to current approaches.

2 Introduction

2.1 Background

Climate change and its effects have developed into a major issue around the world. As the planet warms due to the greenhouse effect, snow melts earlier, hotter weather becomes more frequent, and the amount of summer rain decreases, creating conditions suitable for wildfire ignition. Compiling around 35 years of data, NASA and U.S Forest Service scientists have seen that the fire season starts earlier and ends later, doubling wildfire activity. Additionally, a separate 21-year analysis of satellite data shows that wildfires have also become larger, more intense, and easier to spread [1]. Furthermore, this trend is self-compounding, as greenhouse gases released by wildfires contribute to the same warming that is driving fire seasons to grow longer and more severe, creating a cycle where frequent fires produce conditions for even more future fires.

While fires are a natural and occasionally beneficial part of ecosystems, their increased and uncontrolled activity worsens air quality, contaminates water, forces evacuations, and costs millions to billions of dollars in damage each year [2]. The causes of wildfires can be sorted into two categories: human activity and environmental conditions. Whether it is unattended campfires, arson, or sparks from infrastructure like powerlines, human activity accounts for 85% of all wildfires in the United States [3]. That said, research shows that most human-made wildfires are correlated with periods of fire-conducive conditions in the environment, such as elevated temperatures, high wind speed, and dry vegetation [4].

In conclusion, even when the fire is caused by humans, whether it develops into a wildfire heavily depends on the surrounding conditions, which is why the environment, rather than the ignition source, is the most useful at predicting wildfire risk.

2.2 Rationale

Wildfires have become both increasingly common and destructive due to rising global temperatures, prolonged droughts, and inconsistent climate conditions. As these forest fires continue to multiply in frequency and intensity, identifying high-risk conditions before ignition occurs is critical for improving emergency response and reducing damages. To do this effectively, understanding what causes ignition is necessary.

The environment is the strongest indicator of ignition because wildfire is rarely the result of a singular condition. Rather, ignition appears through the interaction of multiple environmental factors, including but not limited to surface temperatures, vegetation health, fuel availability, forest structures, humidity, and wind conditions [1]. Such conditions are constantly fluctuating, so no single measurement could accurately predict wildfire risk. However, evaluating these environmental indicators together would provide a more ‘complete picture’ of how these factors contribute to combustion, allowing us to understand on a deeper level how wildfire develops in the first place.

Several methods are used to measure these indicators. For example, weather stations continuously measure atmospheric conditions such as temperature, humidity, wind speed, and precipitation. Satellite imagery is used to monitor changes in vegetation, drought conditions, and land cover across a variety of geographic areas. Additionally, reports from the public or emergency responders often serve as added references to speculations of a wildfire occurring [5].

2.3 Past Research

Over the past few years, researchers have extensively explored the topic of CNNs to improve image recognition, object detection, and spatial feature extraction. Many studies have adapted these architectures to analyze UAV imagery and satellite data, demonstrating that CNNs can effectively recognize complex patterns that other methods miss. This body of research has proven CNNs to be a reliable foundation for hazard detection, motivating this framework’s use of dedicated CNNs for each spatial data source: multispectral satellite, multispectral LiDAR, and thermal imagery.

2.3.1 Feature Learning and Automated Spatial Representation

One of the most significant breakthroughs in CNN research was its ability to automatically learn spatial patterns instead of relying on manually engineered descriptions.

Earlier computer vision models required people to decide which visual characteristics mattered, while CNNs progressively learn patterns and combine them into a hierarchical representation. This automatic feature learning allows models to recognize vegetation textures, canopy structures, and anomalous thermal signatures directly from imagery. Most importantly, this framework's thermal, satellite, and LiDAR CNNs rely on this capability to pinpoint environmental patterns without hand-coded thresholds [6].

2.3.2 Mastering Deep Network Optimization

As researchers continued improving these architectures, they discovered that making networks deeper didn't always lead to better performance overall. In many cases, adding more layers made the models harder to train, limiting how deep models could analyze complex spatial information. To address this issue, scholars introduced Residual Networks (ResNets), which use shortcut connections to allow information to bypass certain layers and upgrade the overall flow of the network. By making it possible to train deeper networks without losing valuable information, residual learning significantly improved model stability and feature extraction. This principle is directly relevant here, as deeper CNNs may be needed to capture the subtle environmental patterns present in thermal, satellite, and LiDAR imagery [7].

2.3.3 Multi-Scale Feature Extraction with FPNs

Another challenge was that many models struggled to recognize features appearing at differing spatial scales. This limitation is especially consequential in remote sensing, where small-scale hotspots and large-scale landscape traits both appear within the same satellite imagery. Feature Pyramid Networks (FPNs) addressed this by constructing a multi-scale feature representation that combines detailed spatial information from earlier layers with the stronger contextual understanding of the deeper layers. This capability is particularly relevant to wildfire ignition prediction, since risk indicators can appear as either a small, localized thermal hotspot or a broad pattern of vegetation stress within the same image tile [8].

2.3.4 Deep Learning Through Model Interpretability

As CNNs became increasingly powerful, researchers began prioritizing not only prediction accuracy but also making model decisions understandable. While deep learning (DL) models could uncover complex recurring patterns within larger datasets, their internal reasoning was often difficult to interpret, raising concerns about reliability. To address this, experts developed Gradient-Weighted Class Activation Mapping (Grad-CAM), which uses gradient information to highlight the regions of an image that most influence a model's output. Applying a similar technique to this framework's ignition risk model could allow responders to see exactly which pixels, such as thermal hotspots or stressed vegetation patches, drove a given risk score, adding a layer of interpretability to the system's predictions [9].

3 Methodology

3.1 Solution Overview

The proposed framework is composed of three stages: data collection, data preprocessing, and ignition risk prediction (**Figure 1**). Data will be collected from two complementary sources. A UAV sensor payload, carrying a micro-weather sensor and a mobile multispectral LiDAR (MSL) system, gathers high-resolution, localized data beneath the forest canopy in near real time, while broader thermal and multispectral imagery will be downloaded from the Landsat satellite through the United States Geological Survey (USGS) archive. Combining these sources allows the framework to pair wide-area satellite context with high-resolution, ground-level detail that satellites alone cannot capture due to the thick forest canopy.

Once data for a target region has been collected, it undergoes a preprocessing stage that removes sensor errors and artifacts, normalizes each variable so no single input is unintentionally weighted more heavily than another, and converts raw sensor output into a standardized format that downstream models can read.

Because the collected data falls into two completely distinct categories, spatial imagery and tabular numerical telemetry, the framework routes each through the neural network best suited to it. The four spatial data sources (thermal imagery, multispectral satellite imagery, and MSL-based canopy structure and understory data) are each processed by their own CNN, which learns visual and structural patterns like vegetation stress, canopy density, and thermal anomalies. Meanwhile, the micro-weather telemetry (wind velocity and direction, ambient temperature, humidity, and air pressure) is processed by an MLP, which is better suited to numerical, tabular inputs. The feature outputs of these networks are then fused into a single representation and passed through a final prediction network, which generates a continuously updating ignition risk score and flags high-risk hotspots for early intervention.

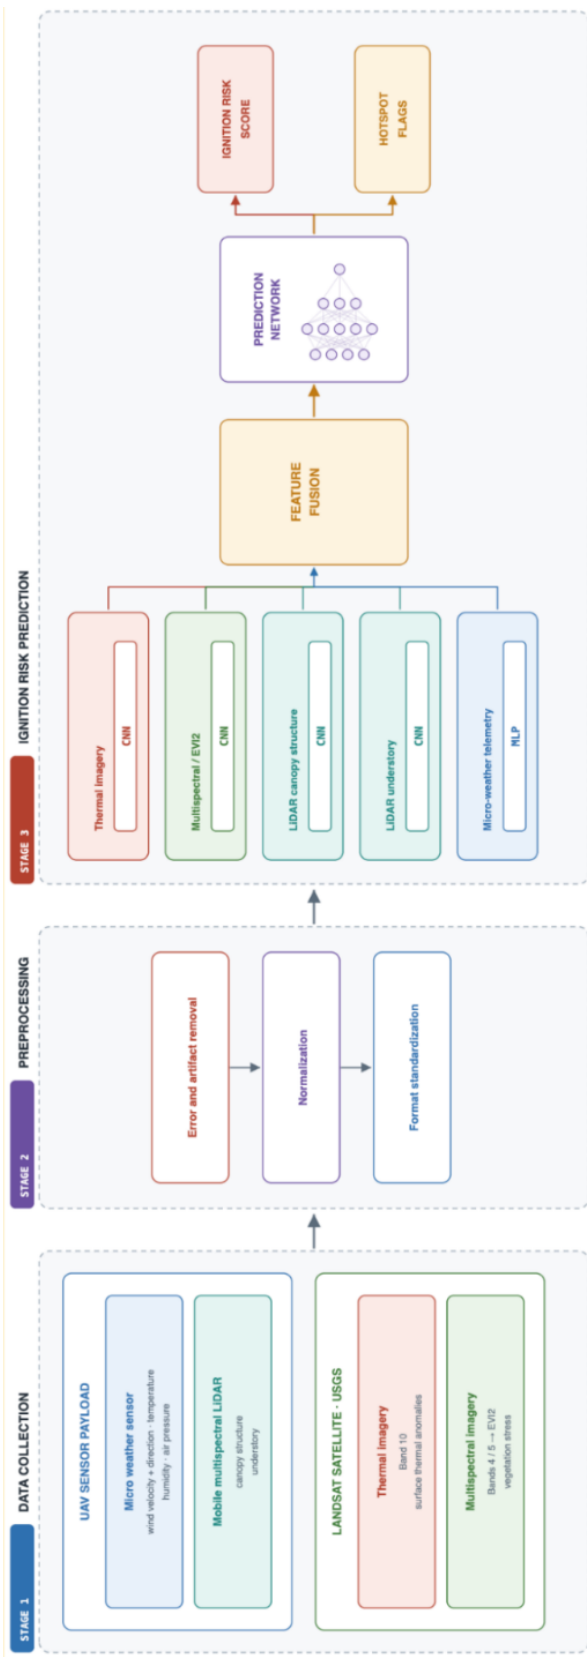

Figure 1: Three-stage architecture of proposed wildfire ignition prediction framework, illustrating data collection, preprocessing, and split CNN/MLP prediction stage feeding a fused prediction network. Created by authors.

3.2 Data Collection

3.2.1 Multispectral Imagery

Multispectral imaging is one of the primary forms of data collection used in our prediction model. “Multispectral” refers to any sensor that collects some number (usually between 3 and 15) of separate wavelength bands across the electromagnetic spectrum [10], and to fulfill our purpose, it will be collected in one of two ways.

3.2.1.1 Multispectral Satellite Imagery

For most use cases, the UAVs will receive downloaded data from the United States Geological Survey (USGS) archive, which is collected by Landsat, a joint program between NASA and the USGS. Landsat satellites play an extremely significant role in measuring and monitoring forest fuel state conditions [11].

The satellites employ an Operational Land Imager (OLI), which collects data in 9 discrete spectral bands simultaneously (hence multispectral), outlined in **Table I**.

Table I: Landsat 8-9 OLI band descriptions. Source: [12].

Band	**Wavelength (μm)**	**Ground Sample Distance (GSD) (m)**
Band 1 – Coastal/Aerosol	0.43-0.45	30
Band 2 – Blue	0.45-0.51	30
Band 3 – Green	0.53-0.59	30
Band 4 – Red	0.64-0.67	30
Band 5 – Near-Infrared (NIR)	0.85-0.88	30
Band 6 – Shortwave-Infrared (SWIR) 1	1.57-1.65	30
Band 7 – Shortwave-Infrared (SWIR) 2	2.11-2.29	30
Band 8 – Panchromatic	0.50-0.68	15
Band 9 – Cirrus	1.36-1.38	30

Of these Bands, NIR and red are key to our data collection, but much conversion needs to be done to the data within the satellite itself before it is ready for preprocessing back on Earth by the CNN [13].

Normally, healthy green leaves, which contain chlorophyll, heavily absorb red light for photosynthesis but reflect NIR light to avoid overheating. However, dry or dying fuel-type leaves do the opposite, reflecting more red light and less NIR [14] (**Figure 2**).

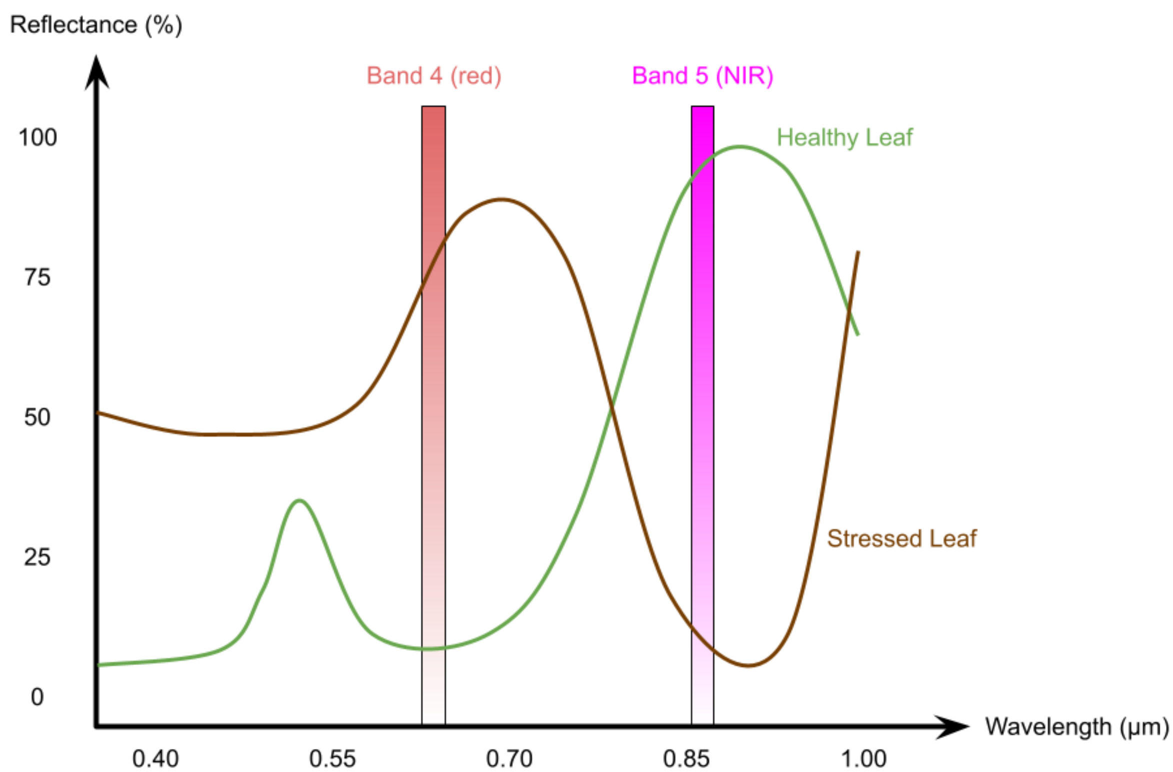

Figure 2: Comparison of spectral reflectance curves for healthy vs. dry/stressed vegetation across Landsat OLI Band 4 (red) and Band 5 (NIR). Created by authors.

To check for this difference, the OLI passes the sunlight through specialized physical filters to isolate the exact wavelengths needed for EVI2. The red light is filtered into Band 4 (0.64-0.67 μm), and the NIR light into Band 5 (0.85-0.88 μm) [13]. The OLI-filtered red and NIR photons then collide with light-sensitive detector crystals. When these collisions occur, electrons are knocked loose, creating a tiny, fluctuating analog electrical current that is directly proportional to the brightness of the given light. The Complementary Metal-Oxide-Semiconductor Read-Out Integrated Circuits (CMOS ROICs) that are bonded to the crystals first amplify the electrical signals with tiny transistors at each pixel site. They then digitize the signals into 14-bit binary through built-in analog-to-digital converters (ADCs), a type of integrated circuit (IC). The CMOS then outputs two separate raw data grids, one for Band 4 (red) and one for

Band 5 (NIR). All of this happens on the satellite, so the data needs to be downloaded from the USGS archive and sent to the UAVs for preprocessing.

3.2.1.2 Multispectral LiDAR

In certain cases, the forest canopy is too dense to be penetrated by NIR and red-light imaging from satellites. Therefore, to provide recent data to the CNN for ignition risk scoring, the UAVs will be equipped with an MSL system [15].

Traditional LiDAR works by emitting a single laser pulse into its surroundings, usually while spinning to map the environment fully. The laser is then reflected off any object and returns to the sensor. After measuring the time it takes for the laser to return to it after emission, the LiDAR calculates the distance between the sensor and the object that the laser bounced off. The distance is measured "using the key components of a lidar system including a GPS that identifies the XYZ location of the light energy and an Inertial Measurement Unit (IMU) that provides the orientation of the plane in the sky (roll, pitch, and yaw)" [16].

However, MSL differs in a few ways that benefit the solution. First, it uses a multitude of lasers, including green, red, and NIR, similar to multispectral satellite imagery, allowing compatibility with EVI2. Furthermore, MSL has significant advantages over traditional monochromatic LiDAR. For example, per the National Library of Medicine, MSL data is much more reliable and accurate in terms of object detection and scene classification, distinguishing a wider range of material textures. When developing three-dimensional (3D) point cloud maps, MSL inputs color and material data directly into each point within the cloud using its multiple laser wavelengths, which can differentiate materials via their spectral signatures. This would allow the UAV to construct a 3D map of the canopy that captures both its structure and the health of the vegetation within it, giving the CNN the spatial and material detail it needs to identify dense, dry, or otherwise high-risk fuel areas that satellite imagery alone cannot resolve beneath thick canopy cover [17].

3.2.2 Thermal Imagery

Thermal imaging is another important type of data that can also be downloaded from the USGS. Onboard Landsat 9 is the Thermal Infrared Sensor (TIRS). According to NASA, TIRS "employs Quantum Well Infrared Photodetectors (QWIPs) to detect… thermal infrared energy" [19]. A QWIP detector is a semiconductor chip that captures electrons, which are then released when thermal radiation of a certain wavelength enters one of the quantum wells. The released electrons "flow to a readout to be recorded and used to create an image of the infrared light source" [20]. TIRS collects 2 different bands of thermal infrared radiation, specifically Bands 10 and 11 (Band 10 will be the only one used in this research, as Band 11 has known calibration issues) [18] (**Table II**). Band 10 collects thermal infrared radiation with wavelengths between 10.60 and 11.19 μm, which lets it determine the proportion of thermal radiation emanating from the surface of the Earth versus the amount of atmospheric interference [19]. By isolating the ground-level thermal signature from atmospheric noise, the system can accurately derive the region's Land Surface Temperature (LST).

Table II: Landsat 8-9 TIRS band descriptions. Source: [12].

Band	**Wavelength (μm)**	**Ground Sample Distance (GSD) (m)**
Band 10 – Thermal-Infrared (TIR) 1	10.60-11.19	100 (resampled to 30)
Band 11 – Thermal-Infrared (TIR) 2	11.50-12.51	100 (resampled to 30)

The ability to measure LST is extremely helpful for estimating forest fuel state, which is the amount of combustible biomass in a forest, such as needles, twigs, and shrubs [20]. The fuel state indicates how easily this vegetation can burn, tracking how dense and dry the materials are in different areas. Plants with sufficient moisture can release enough water to cool their leaves to match the air temperature; however, those with insufficient moisture (indicating dry leaves that contribute to the fuel state) cannot cool quickly enough, causing their leaves to heat [21]. Therefore, TIRS will be able to differentiate between areas of healthy vegetation and areas with stressed, dried, or dead vegetation that are extremely flammable (**Figure 3**). This raw Band-10 signal is later normalized to a 0-1 scale and fed into a dedicated thermal CNN, described further in the 3.3.1 (Preprocessing) and 3.3.2 (CNN) sections, where it is used to detect surface heat anomalies that may point to elevated ignition risk.

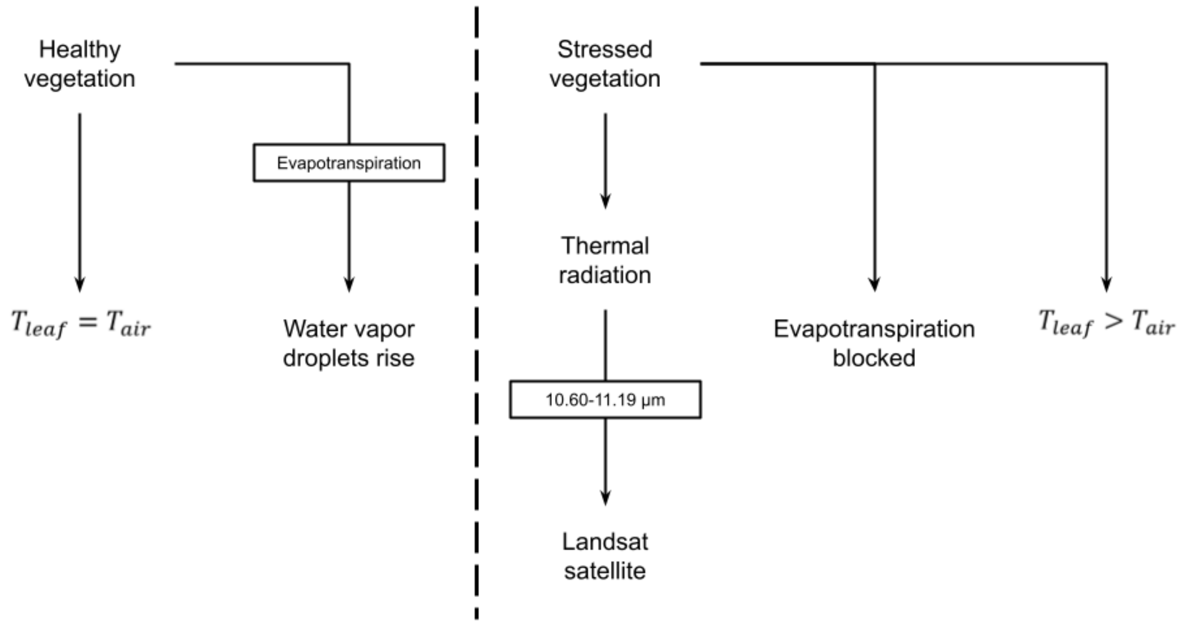

Figure 3: Thermal response mechanism showing transpirational cooling in moisture-sufficient trees versus heat accumulation in moisture-stressed trees detected by Landsat TIRS Band 10. Created by authors.

3.2.3 Micro-weather Sensor Array

The second component of the UAV sensor payload is the micro-weather sensor array, which collects the local wind speed and direction, ambient temperature and humidity, and air pressure. With these weather conditions changing often, the data needs to be collected in near real-time, as they are an important determinate of wildfire ignition risk.

3.2.3.1 Local Wind Velocity and Direction

To collect wind speed and direction, ultrasonic sensors will be implemented. An ultrasonic sensor fires pulses of sound undetectable by humans through the air from one transducer to another. Next, it fires a pulse in the opposite direction (two pairs are covering

North-South and East-West), and compares the time taken for the pulses to reach the receivers. If wind is present, one of the pulses will take less time than the other. The sensor then takes the difference in travel time to calculate the total wind speed $speed_{total}$ by the following equation:

$$speed_{total} = \sqrt{(NS)^2 + (EW)^2}, \tag{1}$$

where NS is the wind speed on the North-South axis, and EW is the wind speed on the East-West axis. To measure the direction of the wind, the sensor uses another formula that includes measurement on both axes [22]:

$$Direction = \arctan 2(NS,\ EW). \tag{2}$$

3.2.3.2 Ambient Temperature

When measuring ambient temperature, the thermal resistor's (thermistor) resistance R changes depending on the temperature. Consequently, the voltage V across the thermistor changes proportionally under a constant current I , according to Ohm's Law:

$$V = IR\ . \tag{3}$$

The sensor then measures the voltage, and the following calibration equation converts the voltage into a temperature reading by the Steinhart-Hart equation:

$$\frac{1}{T} = A + B\ln(R) + C(\ln(R))^3. \tag{4}$$

Here, T is the temperature in Kelvin, R is measured in ohms, and A, B, and C are coefficients provided by the manufacturer [23].

3.2.3.3 Humidity

To measure surrounding humidity, water vapor enters the capacitive humidity sensor's water-absorbing polymer, which sits between two metal plates. As the polymer absorbs moisture, its ability to store electrical charges changes, altering the sensor's capacitance. Electronics within the sensor measure this change and output relative humidity (RH%) through a capacitance-to-digital converter (CDC), another type of IC that determines capacitance by measuring charge time. A calibration equation stored in memory is then used to calculate the RH% from the capacitance [24].

3.2.3.4 Air Pressure

Air pressure is measured using a Microelectromechanical Systems (MEMS) barometer. As air pressure acts on the sensor's silicon diaphragm, it bends slightly and creates stress in the silicon. This stress changes the electrical resistance of the piezoresistors on the sensor, which is then measured and converted into a reading [25].

3.3 Prediction Model

3.3.1 Preprocessing

3.3.1.1 Multispectral Satellite Imagery

By downloading Landsat Collection 2 Level-2 data from the USGS, imagery, that has already been radiometrically calibrated and atmospherically corrected, is received, converting the raw digital numbers into surface reflectance values. Once the multispectral band data has been scaled, the EVI2 can be calculated [26]. Because raw satellite band data cannot be processed easily by the CNN, EVI2 can be used as a preprocessing step to convert data into the correct format. EVI2 is a Vegetation Index (VI) that computes reflectance ratios of NIR (band 5) and red (band 4) light using surface reflectance values to infer the characteristics of vegetation [27].

Compared with other VIs, such as the Normalized Difference Vegetation Index (NDVI) and the traditional three-band Enhanced Vegetation Index (EVI), EVI2 offers extremely low atmospheric vulnerability and reduced sensitivity to background noise while using only two bands, allowing for greater compatibility with both active satellites and neural networks [28].

To compute EVI2, the bands' raw brightness values must first be converted into surface reflectance to remove interference from glare, shadows, and other effects of solar radiation.

Through radiometric calibration, the sensor converts digital numbers (DNs) into radiance and then surface reflectance with atmospheric correction [29]. EVI2 then inputs the NIR and red reflectance ratios for each pixel in the band, as per the formula:

$$EVI2 = 2.5\left(\frac{\rho_{NIR}-\rho_{RED}}{\rho_{NIR}+2.4(\rho_{RED})+1}\right), \tag{5}$$

where ρ_{NIR} is the fraction of NIR light reflected and ρ_{RED} is the fraction of red light reflected [28]. Results from the formula range from –0.74 to 1.25, where numbers closer to 1.25 indicate healthy and dense vegetation, while numbers near 0 denote soil, sparse vegetation, or dead plants. Negative numbers represent water or concrete [30]. After, the data must be split into the appropriate sizes to be inputted into the CNN, which can be done with simple formatting.

3.3.1.2 Multispectral LiDAR

In addition to the multispectral satellite data, the onboard MSL requires preprocessing. Because MSL scans the ground in overlapping flight strips, unadjusted point clouds often contain alignment errors where the same physical feature, like a rock corner or bush shape, appears in two separate locations across adjacent strips. To resolve this, algorithms like Speeded Up Robust Feature (SURF) detect key landmark points in the overlap regions to determine if they represent the exact same physical feature across strips. Once identified, these matching points allow software to align the flight strips; however, candidate matches must first undergo three filtering steps to eliminate false-positive matches: minimizing the distance between SURF feature vectors, using high cross-correlation (CC) to find probable matches, and reflecting geometric characteristics of a matching pattern [31].

When SURF minimizes the distance between feature vectors, it compares them using the Euclidean distance d, which calculates the square root of the sum of the squared differences across each coordinate axis for two feature vectors, $A = (A_1, A_2, \ldots, A_n)$ and $B = (B_1, B_2, \ldots, B_n)$:

$$d(A,\ B) = \sqrt{\sum_{i=1}^{n}(A_i - B_i)^2}, \tag{6}$$

where A_i is the ith value in the first feature vector, B_i is the ith value in the second feature vector, and n is the total dimension of the feature vector. Feature vectors with the smallest distance between each other are noted as probable matches [32]. Then, the CC algorithm compares the intensity and elevation values around the features (collected from the sensor) to determine whether the two features are the same [33]. If a candidate feature passes the first two filters, geometric consistency is checked using Random Sample Consensus (RANSAC). RANSAC finds the translations, rotations, and other transformation estimates from a small set of features. If a certain transformation aligns the features in the smaller set, RANSAC will apply that transformation to the remaining features to see whether they align. If most of the features align, those that do not match after this are noted as outliers and are not used [34].

The next preprocessing step is georeferencing, assigning real-world coordinates to each point in the cloud. The MSL performs much of this work as it maps the point cloud; however, accurate positioning depends on Post-Processed Kinematic (PPK) processing, since live connections can drop out under dense tree canopies or in remote areas [35]. Even with PPK, there are georeferencing adjustments that the drone cannot self-perform during its flight. For example, the physical mounting angles of the MSL and IMU are never perfectly parallel, creating little alignment issues, called boresight angles, that shift the point cloud coordinates [36]. To fix this and other issues, a post-flight strip adjustment method, using the Minimum Hausdorff Distance (MHD), must process the MSL data to calculate the exact angle offsets and prevent potentially overlapping scan lines from appearing blurry or disconnected [37]. Following this, noise filtering removes other measurement errors not already filtered from the MSL by "performing a statistical analysis on each point's neighborhood and trimming those which do not meet a certain criterion" [38]. Finally, as with multispectral satellite images, the sensor uses radiometric calibration to calculate reflectance values, subsequently used to compute EVI2 for each pixel. The next preprocessing step is georeferencing, assigning real-world coordinates to each point in the cloud. The MSL performs much of this work as it maps the point cloud;

however, accurate positioning depends on Post-Processed Kinematic (PPK) processing, since live connections can drop out under dense tree canopies or in remote areas [35]. Even with PPK, there are georeferencing adjustments that the drone cannot self-perform during its flight. For example, the physical mounting angles of the MSL and IMU are never perfectly parallel, creating little alignment issues, called boresight angles, that shift the point cloud coordinates [36]. To fix this and other issues, a post-flight strip adjustment method, using the Minimum Hausdorff Distance (MHD), must process the MSL data to calculate the exact angle offsets and prevent potentially overlapping scan lines from appearing blurry or disconnected [37]. Following this, noise filtering removes other measurement errors not already filtered from the MSL by "performing a statistical analysis on each point's neighborhood and trimming those which do not meet a certain criterion" [38]. Finally, as with multispectral satellite images, the sensor uses radiometric calibration to calculate reflectance values, subsequently used to compute EVI2 for each pixel.

To extract canopy height, the raw 3D multispectral point clouds are converted into two-dimensional (2D) raster channels [39], generating a Digital Terrain Model (DTM) and a Digital Surface Model (DSM). Subtracting the DTM from the DSM provides the Canopy Height Model (CHM). These height features are then stacked with calculated data from EVI2 to transform the data stream into a multi-channel 2D image tensor, fed to the CNN [40]. Generating this CHM is necessary because tree height significantly affects potential ignition and fire spread. The taller a tree gets, the larger its radiant heat flux, drastically increasing the chances of igniting surrounding vegetation and spreading the fire [41].

Preprocessing to acquire Canopy Bulk Density (CBD) requires extracting a larger depth of 3D complexity from the MSL before being fed to the CNN. The point clouds then undergo voxelization to convert 3D spaces into a uniform grid [42], where each voxel (volumetric pixel) contains coordinate geometries and multispectral intensity metrics [43]. These voxels are promptly fed into the CNN, enabling the network to more accurately understand vegetation distribution and volumetric density, which is important because CBD, like canopy height, is a determining factor in fire behavior and ignition risk scoring. In fact, according to the National Wildfire Coordinating Group (NWCG), "CBD estimates are used to determine the threshold spread rate, or surface wind speed, used to determine the likelihood of active crown fire" [44].

Lastly, capturing understory vegetation data demands methods that leverage MSL's ability to penetrate the top canopy layers. The MSL returns are heavily filtered to isolate below-canopy structural strata amid dense tree crowns and ground-level noise. These isolated 3D point clouds are decomposed and projected onto 2D multispectral imagery, creating a dataset of understory geometry. The CNN, with this dataset, can reliably interpret the understory vegetation, which is essential for two reasons. First, it is a critical indicator of forest health [45], and second, more importantly, it will allow the CNN to identify potential ladder fuels, certain understory flora that that provide an upward path for fire to travel to taller trees [46], resulting in crown ignition and fire spread (as mentioned previously).

To allow the CNN to optimize training across all extracted MSL features (canopy height, CBD, and understory vegetation), a final normalization step is required. Because these metrics originate from different physical units and scales (such as meters for height versus point counts for voxel density), feeding raw values directly into the network would cause features with larger numerical magnitudes to dominate model training. To prevent feature bias, Min-Max Normalization is applied to scale all feature values into a uniform range between 0 and 1 using the equation:

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}}, \quad (7)$$

where X represents the raw feature value (e.g., height, voxel density, or understory point density), X_{min} is the minimum value observed for that feature across the dataset, and X_{max} is the maximum value. Standardizing all MSL-derived metrics to this 0-to-1 range accelerates CNN training convergence and ensures that each structural driver of ignition risk is weighted proportionally by the network [47].

3.3.1.3 Thermal Imagery

After downloading the thermal Band 10 imaging from the Landsat database, much of the preprocessing, namely calibration and geometric correction, has already been done by USGS (similar to the multispectral data streams) [26], leaving only a few further steps before the data can be networked.

Because clouds, their shadows, and defective sensor pixels can introduce inaccuracies in the data, they must be removed. USGS provides an algorithm (CFMask) to acknowledge and accurately “flag” these potential problems by using “clear view forest as a reference to define cloud boundaries for separating cloud from clear view surfaces in a spectral-temperature space. Shadow locations are predicted according to cloud height estimates and sun illumination geometry, and actual shadow pixels are identified by searching the darkest pixels surrounding the predicted shadow locations” [48]. Additionally, satellite images cover broad areas, hence they need to be split and formatted into smaller patches, each matching in size to the EVI2-calculated multispectral image bands. To allow for use as input, once again, Min-Max Normalization is implemented to scale each pixel value in the patches.

3.3.1.4 Micro-Weather Sensor Array

When preparing the drone-collected micro-weather data for the MLP, the raw data must undergo important preprocessing steps before it is fed to the neural network. First, the data must be cleaned, through handling (either removing or imputing) missing or impossible values and adjusting noisy data [49]. Specific to wind telemetry, the wind direction calculated from Formula (2) outlined in 3.2.3.1 (Local Wind Velocity and Direction) is output in degrees of azimuth (0 to 360 degrees). An MLP cannot easily process this format; for example, it may think that 0 and 360 degrees are on opposite sides, when, in reality, they are in the same direction (explained further in **Table III**). To fix this problem, the wind speed and direction must be converted into two linear Cartesian components: u-wind, which represents the east-west (zonal) component, and v-wind, which represents the north-south (meridional) component, where positive values indicate wind blowing east and north, respectively, and vice versa for negative values, with the magnitudes of these values indicating speed [50]. These conversions can be calculated as per the following equations:

$$u = -speed_{total}(\sin(\theta)) \tag{8}$$

and

$$v = -speed_{total}(\cos(\theta)), \tag{9}$$

where $speed_{total}$ again represents the wind speed and θ is the angle of the wind direction in degrees, greatly simplifying the numbers and increasing continuity for the MLP [51].

Table III: Continuity advantage of Cartesian vectors compared to raw degrees. Created by authors.

Raw Degrees	**Cartesian Vectors**
Discontinuous • 0° → Scalar: 0 • 360° → Scalar: 360 • Δ = 360 (Appears far)	Continuous • 0° → ($u = 0$, $v = -speed_{total}$) • 360° → ($u = 0$, $v = -speed_{total}$) • Δ = 0 (Identical point)

Furthermore, because different micro-weather sensors record data at varying sampling frequencies, data synchronization is required to ensure that all measurements correspond to the exact same time step. One effective approach is timestamp matching, where a system clock assigns timestamps to each reading, and any data falling within a predetermined window is grouped together to reflect conditions at that designated moment [52]. If a sensor fails to record a reading during collection, a null value is generated in the dataset. Because an MLP cannot process null entries, these missing values are replaced with the last valid reading through a process known as forward fill [53]. Moreover, environmental turbulence or sensor noise can create artificial spikes that distort true ambient conditions. To eliminate these errors without leaving gaps in the dataset, a moving average filter is applied to calculate the mean of surrounding values within a small-time window and replace anomalous readings. Finally, Min-Max Normalization is used once more to scale all variables into a 0-to-1 range, ensuring that each micro-weather factor is weighed equally in the MLP's predictions.

3.3.2 CNN for Spatial and Structural Analysis

CNNs, a deep learning architecture designed for visual pattern recognition, are used to analyze the high-resolution environmental imagery collected by the UAV and satellite sensor systems [54]. While the MLP focuses on numerical micro-weather data, the CNN specializes in processing the spatial data prepared in the preprocessing stage, allowing the model to interpret what the landscape looks like rather than relying only on single numerical values. By recognizing patterns in vegetation, canopy structure, and heat signatures, this predictive prototype builds a deeper understanding of the environmental conditions that contribute to wildfire ignition and spread.

Unlike traditional methods that require researchers to identify specific environmental indicators manually, a CNN automatically learns which spatial features are most relevant by repeatedly training on historical, labeled environmental imagery where the resulting wildfire outcome (ignition or no ignition) is already known. During this training process, the CNN's internal filters are adjusted through backpropagation, which calculates how each weight contributed to the model's error, and gradient descent, which uses that calculation to update the weights [55], so that the patterns most associated with real ignition events are weighted more heavily. This allows the network to uncover relationships between landscape features that may not be immediately visible to human researchers.

3.3.2.1 Input Standardization and Training

Because the CNN can only process data in a consistent numerical format, every input source described in preprocessing is converted into stacked 256×256-pixel image tensors [56] before being fed into the network. This tile size lets the CNN evaluate localized terrain conditions while still being small enough to individually compare them consistently across different geographic regions, allowing the model to identify areas where multiple wildfire risk factors overlap. Each channel of the tensor is scaled differently depending on how the CNN is meant to interpret it. The EVI2 channel is left on its naturally computed scale of -0.74 to 1.25, since this range already carries meaning established in 3.3.1.1 (Multispectral Satellite Imagery) [30]. The thermal channel and the MSL-derived data channels are each scaled from 0 to 1 using Min-Max Normalization [47], matching the normalization already applied to that data in preprocessing. This input standardization is illustrated in **Figure 4** below.

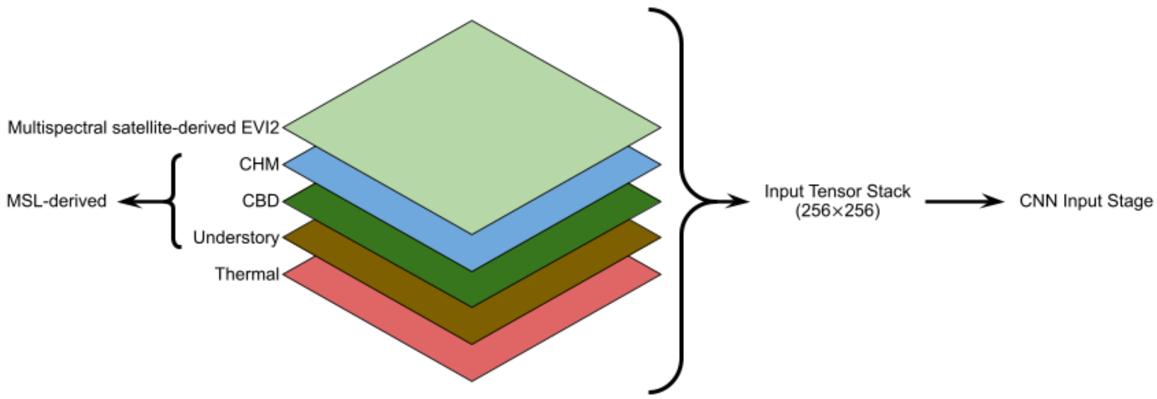

Figure 4: Schematic of the 256×256 multi-channel input tensor stack combining aligned spatial channels (EVI2, CHM, CBD, understory, and thermal layers) prior to CNN ingestion. Created by authors.

During training, the CNN is presented with many of these stacked tensors, along with the known wildfire outcome for that location and time. Through repeated exposure, the network's convolutional filters learn which combinations of EVI2 values, temperature values, and canopy height values are statistically associated with ignition, allowing the trained model to recognize similar combinations in new, unseen data.

3.3.2.2 Vegetation Health Analysis (EVI2)

As calculated in 3.3.1.1 (Multispectral Satellite Imagery), EVI2 is derived from multispectral satellite imagery [27] already provides a value for every pixel in a scene. Rather than reading these EVI2 pixel values in isolation, the CNN evaluates them as part of a spatial neighborhood, allowing it to recognize repeating patterns across the tile that are associated with moisture stress, vegetation density, and fuel continuity [57]. Because the CNN has been trained on historical EVI2 data paired with known ignition outcomes, it has learned which arrangements of EVI2 values, not just individual values, correspond to drier, more flammable vegetation. This allows the model to interpret the EVI2 channel as a continuous fuel-availability signal rather than a simple health score for each isolated pixel.

3.3.2.3 Multispectral LiDAR (Canopy Structure)

As outlined in 3.3.1.2 (Multispectral LiDAR), the CHM is derived by subtracting the DTM from the DSM obtained via the UAV-mounted MSL point cloud [58]. The CNN does not create this model itself; instead, it takes the already-normalized canopy height raster as one stacked input channel, alongside the voxelized CBD data and the projected understory point clouds, also described in 3.3.1.2 (Multispectral LiDAR).

From these structural channels, the CNN learns to recognize the 3D arrangement of vegetation throughout the tile, making it possible to detect ladder fuels, layers of vegetation that let fire climb from the ground into the tree canopy [46]. Because the LiDAR channel is normalized to the same 0-1 scale as the thermal channel, the CNN weighs structural risk equally alongside temperature-based risk when learning ignition patterns [59].

3.3.2.4 Thermal Imagery

The thermal channel is built from Landsat Band 10 imagery, which, as explained in 3.3.1.3 (Thermal Imagery), is cloud-masked, patch-split, and normalized to a 0-1 scale before being stacked into the CNN's input tensor. Because Landsat 8 and Landsat 9 operate together, a given location is only revisited approximately every 8 days [26], so the CNN is not evaluating a continuous live feed; rather, it is trained to recognize abnormal heat signatures within each available thermal snapshot, relative to what a normal, non-fire temperature pattern looks like for that terrain.

By training on normalized temperature values across many pixels rather than a single reading, the CNN learns to recognize temperature gradients and localized clusters of unusually high values relative to their surroundings, rather than reacting only to raw heat magnitude [60].

This enables the network to flag developing thermal anomalies that may indicate increased ignition potential before they are visible through traditional observation.

3.3.2.5 Spatial Feature Extraction

Instead of allowing each preprocessed channel to independently predict wildfire risk, the CNN extracts spatial features jointly across the entire stacked tensor for each geographic tile [61]. This prevents isolated, single-source predictions and produces a unified representation of the landscape. Through multiple convolutional layers, the CNN pools the EVI2 values, normalized temperature values, and normalized canopy height values in each tile into compact feature representations that capture the thermal, structural, and ecological characteristics associated with wildfire risk. Earlier convolutional layers identify simple patterns such as edges, textures, and localized temperature variation, while deeper layers combine these details into more complex relationships between vegetation structure and surface heat. This layered approach allows the CNN to move beyond any single data point and understand how conditions across a tile interact to create favorable conditions for ignition.

3.3.2.6 Outputs

The CNN's role is to convert the preprocessed EVI2, thermal, and canopy structure channels into a compact set of learned spatial features for each tile rather than producing a standalone wildfire prediction. These flattened spatial features are combined with the meteorological features generated by the MLP, allowing the larger framework to evaluate spatial and atmospheric risk factors together instead of separately. This combined representation ultimately drives the framework's continuous wildfire ignition probability score for each geographic location, which is later converted into the spatial risk heatmaps and hotspot alerts described in the ignition risk scoring section.

3.3.3 MLP for Micro-weather Analysis

Within the proposed analytical framework, an MLP, a class of predictive artificial neural networks, is implemented as a specialized engine for processing abstract, heterogeneous meteorological telemetry. While spatial data such as multispectral and thermal imagery is routed through CNNs to recognize patterns, the MLP is refined for high-resolution, structured data streams from UAV micro-weather sensor arrays.

3.3.3.1 Input Parameters and Data Streams

The MLP architecture processes micro-weather data collected beneath the canopy in real time, which provides higher quality intelligence than broad regional reports from data sources like the National Weather Surface (NWS) [62]. The primary inputs for the MLP include local wind velocity and direction, ambient temperature and humidity, and air pressure data, outlined below and in **Figure 5**.

3.3.3.1.1 Local Wind Velocity and Direction

Local wind velocity measures the speed of airflow, while wind direction indicates how air moves beneath the forest canopy, providing detailed information about localized atmospheric conditions. This data would be used to assess conditions that influence the dryness of surrounding vegetation and identify environments prone to ignition. Higher wind speeds accelerate moisture evaporation from vegetation and increase oxygen availability, causing grasses, shrubs, and leaf litter to dry more rapidly and become more susceptible to ignition. The MLP processes wind velocity and direction together because these variables describe local airflow patterns that influence fuel drying and wildfire ignition potential.

3.3.3.1.2 Ambient Temperature and Humidity

Ambient temperatures measure the temperature of the surrounding air at the UAV's location, while relative humidity measures the amount of water vapor present in the atmosphere.

Higher ambient temperatures increase the rate at which moisture evaporates from vegetation, causing leaves, grass, and other fuels to dry more quickly.

These variables would directly influence the Live Fuel Moisture (LFM) index and the drying rate of forest fuels, which would serve as primary indicators of ignition probability. The MLP processes both variables together because they are highly correlated; for example, high temperatures combined with low relative humidity generally indicate greater wildfire ignition potential than either variable alone.

3.3.3.1.3 Air Pressure

Localized sensors integrated into UAV sensor arrays would provide high-frequency data to distinguish the atmospheric stability of micro-weather zones. The MLP would use these readings to identify specific pressure gradients and barometric trends that often precede sudden shifts in local wind velocity or abrupt drops in relative humidity.

Due to the live-feed architecture, air pressure would serve as an important variable for identifying unstable environmental states. Specific pressure patterns can indicate low-level air mass movement, which would intensify the drying of fine fuels.

In addition, this barometric data is used by the network to cross-reference and normalize other meteorological signals, ensuring that weighted feature vectors accurately represent the atmospheric risk state before they are fused with fuel state imagery.

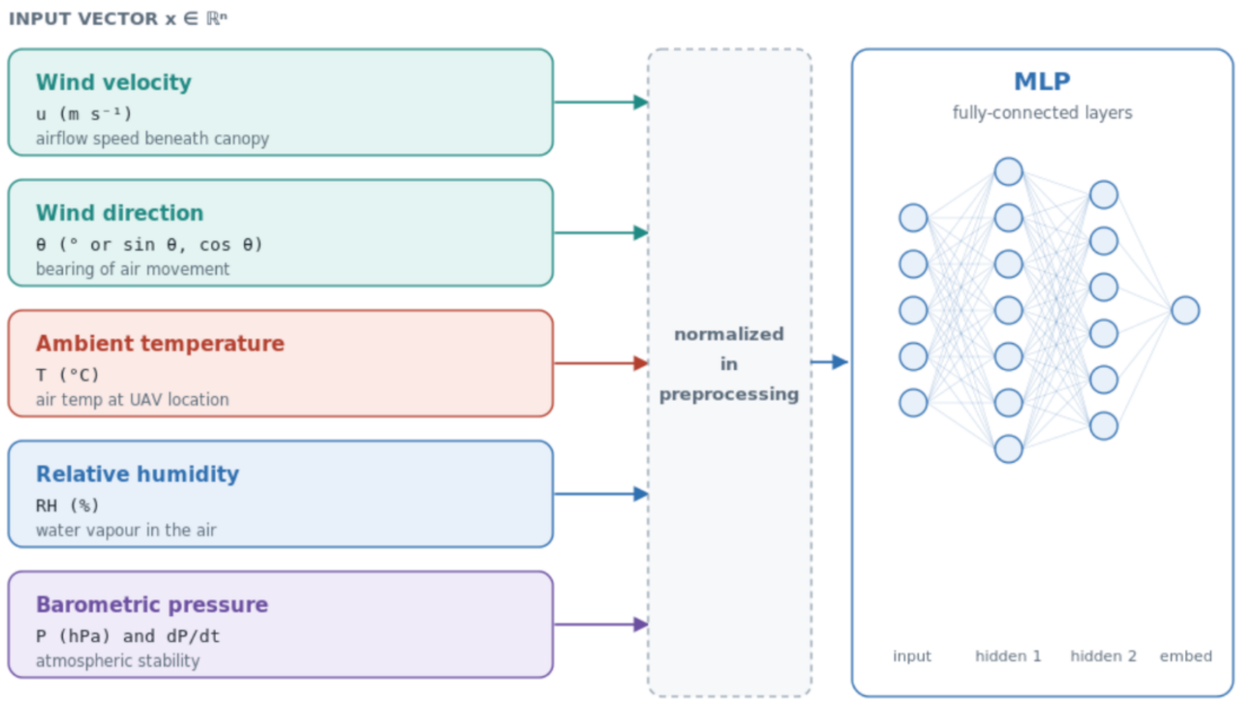

Figure 5: Micro-weather inputs to MLP, collected beneath canopy and passed as a structured feature vector to network's dense layers. Created by authors.

3.3.3.2 Model Training and Label Preparation

Before the MLP can generate a meaningful risk output, it must first be trained on a labeled dataset consisting of paired historical inputs and known outputs. Creating this training set would require two components: a source of environmental input data that matches the MLP's parameters, and a source of ground-truth labels indicating whether ignition occurred under those conditions.

3.3.3.2.1 Ground Truth Labels

Historical wildfire occurrence records, such as those maintained by databases like the Fire Program Analysis Fire-Occurrence Database (FPA FOD) or satellite-based active fire detection products (e.g., MODIS or VIIRS), provide the location and timestamp of confirmed past ignitions [63].

Each ignition event serves as a confirmed training example (labeled as 1). Because a model trained only on positive examples cannot learn to differentiate risky conditions from safe ones, a proportionate number of negative examples (labeled as 0) must also be sampled from the same geographic regions in time periods where environmental conditions were similar, but no ignition occurred.

3.3.3.2.2 Input-Output Pairing

For each labeled positive or negative event, the corresponding environmental conditions—wind velocity and direction, ambient temperature, relative humidity, and air pressure—must be reconstructed for that specific time and location. Because UAV micro weather sensors were not deployed historically, historical reanalysis or ground station weather data (e.g., NOAA or ERA5 records) would be used as a proxy for what a UAV sensor would have measured at that coordinate [64].

Each input-output pair therefore consists of a feature vector, including wind, temperature, humidity, and pressure, matched to a binary ignition label for that location and time window (**Table IV**).

Table IV: Example of model training using illustrative values. Created by authors.

Assembled Supervised Dataset							
$u\ (\frac{m}{s})$	$\sin(\theta), \cos(\theta)$	T (°C)	RH (%)	$P\ (hPa)$	$\frac{dP}{dt}$	$\rightarrow y$	$x \in \mathbb{R} \mapsto y \in \{0, 1\}$
2.4	0.71, 0.70	34.2	18	1008.1	-1.4	→ 1	• One row = one location/one time window • This is the unit the MLP is trained on
0.8	-0.26, 0.97	21.6	62	1015.3	+0.2	→ 0	
3.9	0.94, 0.34	37.8	12	1005.7	-2.1	→ 1	

3.3.3.2.3 Train/Validation/Test Split

The assembled dataset is divided into training, validation, and test sets. To prevent data leakage and provide a realistic estimate of real-world performance, the split should be based on time period or geographic region rather than random sampling. For example, the model can be trained on fire seasons from earlier years and validated on a held-out, more recent season, ensuring it is not evaluated on conditions it has already seen.

3.3.3.2.4 Loss Function

Since the model's task is to distinguish ignition-prone conditions from non-ignition conditions, the MLP is trained using a binary cross-entropy loss function, which penalizes the model based on how far its predicted ignition probability deviated from the true label at each training example.

3.3.3.3 Processing Logic

The processing logic of the proposed framework is centered on a split DL architecture designed to turn heterogeneous data into predictive ignition intelligence. This stage reconciles high-resolution, high-frequency UAV telemetry with broad, multispectral satellite baselines through a three-step refinement process.

3.3.3.3.1 Coordinate Projection

The system performs a Universal Transverse Mercator transformation on all incoming GPS data. This converts latitude and longitude into metric-based northing and easting coordinates, which is essential for the precise alignment of hotspot data with vegetation indices [65].

3.3.3.3.2 Atmospheric Calibration

Raw digital numbers from Landsat Band 4 (red), Band 5 (NIR), and Band 10 (TIRS 1) are calibrated into surface reflectance and brightness temperature values, detailed in 3.3.1.1 (Multispectral Satellite Imagery). This step removes disturbance from solar radiation and atmospheric haze to ensure that the EVI2 calculation is accurate [26].

3.3.3.3.3 Signal Filtering

Particulate matter and gas concentration telemetry from UAV sensors are filtered to extract chemical signatures of combustion. This process will eliminate environmental background noise to identify potential ignitions before they are visually apparent [66].

3.3.3.4 Feature Extraction

3.3.3.4.1 Nonlinear Relationship Mapping

The MLP uses dense hidden layers to identify combined risk patterns among atmospheric variables. For example, it extracted the hidden correlation between a rapid drop in barometric pressure and a corresponding spike in localized temperature to identify volatile environmental states [67].

3.3.3.4.2 Micro-weather State Vectoring

The output of the MLP is a weighted feature vector that represents the immediate atmospheric risk state of a specific coordinate, ensuring that temporary fluctuations in LFM and wind velocity are mathematically represented before data fusion [68].

3.3.4 Ignition Risk Scoring

3.3.4.1 Cross-Modal Feature Fusion

The primary intelligence of the engine is the feature fusion stage, where the spatial outputs of the CNN are combined with the tabular outputs of the MLP.

3.3.4.1.1 Hybrid Architecture Integration

High-dimensional feature maps representing the fuel state (biomass density and EVI2 health) are flattened and concatenated with the MLP’s meteorological feature vectors. This ensures the model accounts for both the fuel's readiness and the volatility of atmospheric conditions [69].

3.3.4.1.2 Probabilistic Ignition Scoring

This fused dataset is processed to output the final Integrated Risk Map, which provides a high-resolution, percentage-based probability of combustion for every coordinate in the target zone. This computational output would allow responders to identify high-risk intervention zones before a wildfire ignites [70].

3.3.4.2 PINN

The Physics-Informed Neural Network (PINN) is implemented and developed to estimate the probability of wildfire ignition by dynamically combining spatial, temporal, structural, and meteorological information extracted throughout the pipeline. Unlike traditional neural networks that rely solely on historical data, the PINN incorporates established wildfire physics into the learning process, ensuring that the resulting predictions are both highly data-driven and physically consistent with real-world dynamics.

The primary objective of the network is not the detection of an active fire, but rather the determination of whether current environmental conditions are sufficiently favorable for wildfire ignition within a 12- to 24-hour window. By identifying these conditions early, the system enables proactive mitigation strategies before any visible fire activity occurs. Further, the PINN incorporates physical constraints directly into the training phase, allowing predictions to remain rigidly consistent with governing wildfire dynamics even when training data is sparse or noisy [71].

3.3.4.3 Input Feature Representation and Architectural Integration

The network processes a fused feature vector generated from all previous stages of the broader architecture. This input layer relies on three distinct computational streams to capture the multi-faceted nature of wildfire risks: spatial relationships (which are handled natively by the PINN), spatial features, and environmental features.

3.3.4.3.1 Spatial Features

Spatial features are extracted via the CNNs from aerial imagery and thermal observations, described in 3.3.2 (CNN for Spatial and Structural Analysis). This stream captures critical visual indicators, including thermal anomalies, a vegetation health index (EVI2), fuel continuity, canopy structure, and surface temperature distributions.

3.3.4.3.2 Environmental Features

Environmental variables are processed using the MLP, by handling numerical meteorological data [72], discussed in 3.3.3 (MLP for Micro-weather Analysis). The variables ingested by this network include fuel moisture levels, ambient air temperature, relative humidity, wind speed and direction, and atmospheric pressure.

3.3.4.4 Physics-Informed Feature Fusion and Optimization

Following initial processing, the outputs from the CNN and MLP are concatenated into a unified, high-dimensional feature representation. Unlike conventional feature fusion techniques that treat variables uniformly, the PINN is optimized through a dual objective strategy that jointly respects observational data and wildfire physics.

During optimization, the network minimizes two loss terms simultaneously. The first, data loss, measures how closely the model's predictions match historical ignition events. The second, physics loss, penalizes predictions that violate fundamental wildfire behavior, including heat transfer, radiant heat flux, wind transport, terrain effects, and fuel-driven fire dynamics. By combining data mismatch with partial differential equations (PDEs) or physics-based residual losses, the learned solutions remain naturally consistent with governing physical equations, including:

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \lambda(\mathcal{L}_{physics}). \tag{10}$$

In this formulation, $\mathcal{L}_{total}$ is the total loss function used to optimize the PINN, $\mathcal{L}_{data}$ represents the empirical prediction error against the ground truth, $\mathcal{L}_{physics}$ penalizes violations of established wildfire physics, and the hyperparameter λ controls the relative influence of the physical constraints during backpropagation [73]. This architecture is illustrated in **Figure 6** below.

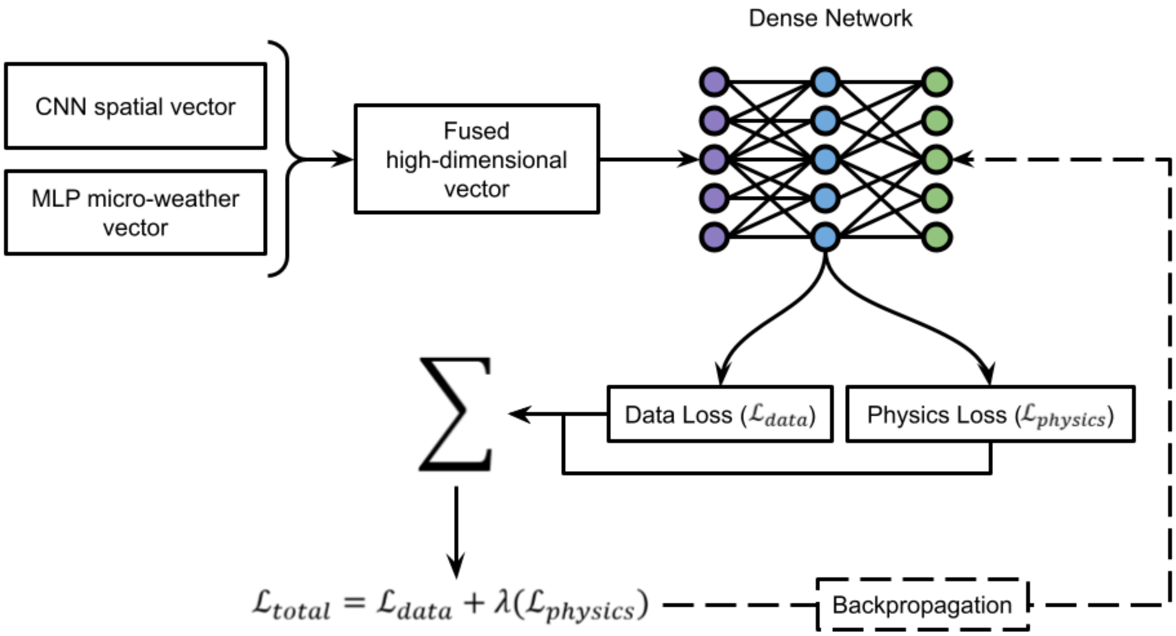

Figure 6: PINN architecture depicting feature fusion (CNN + MLP) and backpropagation via a dual loss function combining empiral data mismatch ($\boldsymbol{\mathcal{L}_{data}}$) and physical PDE residual penalties ($\boldsymbol{\mathcal{L}_{physics}}$). Created by authors.

3.3.4.5 Predictive Modeling and Ignition Risk Output

The fused environmental representation is passed into the core PINN layer, which learns the complex nonlinear interactions among vegetation, terrain, meteorology, and evolving temporal environmental trends. Rather than forcing a binary decision, the model outputs a continuous probability score to reflect the nuanced gradients of environmental risk. The final risk prediction is computed at the output layer using a Sigmoid activation function:

$$P(ignition) = \frac{1}{1+e^{-z}}. \tag{11}$$

Here, z represents the raw network output vector, and the final probability $P(ignition)$ is bound within the interval [0,1]. This continuous score represents the estimated likelihood that local environmental conditions are favorable for wildfire ignition within the target forecast window.

3.3.4.6 Model Outputs and Operational Interpretability

The PINN pipeline delivers four primary actionable outputs designed to interface directly with emergency management systems. First is ignition probability, a continuous risk score indicating the likelihood and severity of ignition. For example, a score of 0.08 represents a low ignition likelihood, 0.41 represents moderate susceptibility, and 0.86 denotes a high likelihood of ignition. Next is risk classification, which translates these continuous scores into operational alert levels; the probabilities are grouped into four distinct categories: low, moderate, high, and extreme. Additionally, the system generates a spatial risk heatmap, projecting risk probabilities geographically to produce high-resolution ignition-risk maps [74]. These web-GIS layers allow emergency managers to visualize developing hotspots across the monitored landscape in real time (**Figure 7**). Finally, confidence estimation provides a secondary metric calculated alongside the ignition probability [75]. This identifies areas of high model uncertainty where additional sensing, field validation, or satellite observations may be required to confirm the risk.

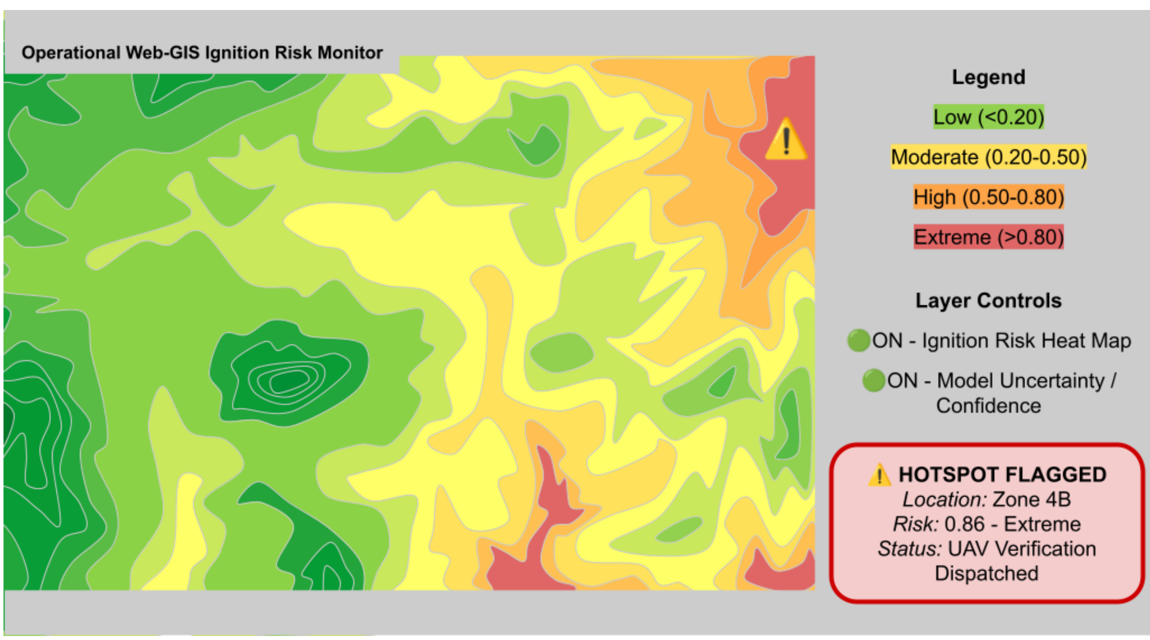

Figure 7: Example operational web-GIS monitoring interface displaying high-resolution spatial ignition probability heatmaps, 4-tier categorical risk classification, model uncertainty/confidence layer toggles, and automated hotspot alerts for targeted UAV verification dispatch. Created by authors.

3.3.4.7 Model Evaluation Metrics

Prediction performance is continuously evaluated through standard machine learning metrics for risk verification. The ROC AUC assesses overall classification performance across decision thresholds. Precision measures the proportion of predicted ignitions that are correct, while Recall measures the network's ability to identify true ignition events. The F1 Score provides a balanced metric that accounts for both false positives and missed ignitions. Finally, a calibration error evaluates whether predicted probabilities accurately reflect real-world ignition frequencies over time.

3.3.4.8 Operational Applications

The physics-informed ignition risk network shifts wildfire management from a reactive framework to a proactive one. Operationally, the system generates real-time ignition probability maps that enable agencies to prioritize UAV deployment in high-risk regions for physical verification. By triggering early hotspot alerts before visible wildfire development occurs, the network supports strategic resource allocation, targeted fuel mitigation planning, and rapid emergency response scheduling.

4 Discussion

The proposed framework offers a more integrated approach to wildfire ignition prediction than conventional monitoring systems provide. Rather than relying on a single sensing modality, it unifies multiple remote sensing platforms, automated data processing pipelines, and specialized neural network architectures into a cohesive system. This design reflects the understanding that wildfire ignition stems from complex interactions among vegetation, terrain, and atmospheric conditions rather than any single variable. Furthermore, it directly addresses the limitations of existing monitoring tools—such as lookout towers and satellite-only systems—which suffer from cloud cover obstruction, coarse spatial resolution, and infrequent revisit times that hinder the detection of subtle pre-ignition changes. By pairing wide-area satellite coverage with high-resolution UAV telemetry collected beneath the canopy, this framework is equipped to detect localized risk factors that satellite-only systems miss entirely.

Central to this advancement is a dual-platform data collection strategy that combines Landsat satellite imagery with UAV-based measurements to capture both broad regional trends and granular, ground-level detail. This pairing yields a significantly more comprehensive assessment of wildfire conditions than either platform could achieve independently. However, expanding environmental visibility introduces operational trade-offs, including constrained UAV flight endurance, battery limitations, vulnerability to adverse weather, and the elevated costs associated with sophisticated sensor payloads.

The preprocessing pipeline serves as the essential bridge between raw sensor data and deep learning analysis, standardizing heterogeneous inputs and correcting environmental noise. Transforming complex three-dimensional (3D) point clouds into structured 2D representations enables efficient processing by the CNNs; however, this dimensional reduction introduces a trade-off between computational efficiency and structural detail. Simplifying 3D geometries into 2D raster formats risks omitting fine-scale canopy and understory features that could influence fire behavior. Additionally, performing real-time atmospheric corrections and signal processing demands substantial computational resources, potentially impacting system responsiveness during field operations.

The system architecture further optimizes environmental assessment by routing distinct data modalities to the neural network models best suited to interpret them. Convolutional Neural Networks (CNNs) extract spatial features—such as vegetation indices, surface textures, and canopy structural metrics—while a Multilayer Perceptron (MLP) processes dynamic atmospheric data from numerical weather telemetry. Combining these models enables simultaneous evaluation of static landscape features and dynamic microclimatic conditions. Nevertheless, this hybrid approach inherits inherent deep learning challenges: these models require extensive, high-quality training datasets to generalize effectively, and their internal decision-making processes remain notoriously opaque. This lack of interpretability presents a hurdle for emergency responders, who must trust and act on automated risk assessments. Explainability techniques such as Grad-CAM offer a viable solution by visualizing the input regions that most heavily influence a model's prediction, serving as a natural next step to enhance the transparency of the framework's CNNs.

Finally, the data fusion stage synthesizes outputs from both the CNNs and the MLP to reveal complex interactions among terrain, vegetation, and microclimates. This integration enables a continuously updated wildfire probability score that supports proactive emergency management. However, combining multiple independent data streams demands rigorous spatiotemporal synchronization between satellite observations and UAV telemetry. Any misalignment in spatial coordinates, timestamping, or sensor calibration across the pipeline could compromise the accuracy of the final risk score, as data fusion fundamentally relies on the assumption that all inputs correspond to the same physical location and time frame.

Ultimately, these trade-offs indicate that while the proposed framework delivers a significantly more proactive and granular approach to wildfire ignition prediction than conventional tools, its real-world efficacy hinges on data quality, precise sensor synchronization, and access to robust training datasets. The following concluding section outlines actionable strategies to validate, refine, and scale this framework in future research.

5 Conclusion

This study demonstrates the potential of combining machine learning with multi-modal remote sensing to address the escalating threat of wildfires by predicting ignition risk before fires occur. By pairing Landsat satellite imagery with high-resolution UAV telemetry and routing spatial inputs—such as thermal, LiDAR, and vegetation imagery—through dedicated CNNs alongside a MLP for real-time meteorological data, this framework transforms complex environmental metrics into continuous ignition probability scores and spatial risk heatmaps. These actionable outputs are designed to support earlier, better-informed decision-making for emergency responders and vulnerable communities. However, realizing this potential requires addressing the critical trade-offs highlighted in this report, particularly regarding spatiotemporal sensor synchronization, training data availability, and model interpretability. Future research should focus on strengthening real-time edge computing capabilities, expanding training datasets to encompass diverse microclimates and forest ecosystems, and integrating advanced architectures—such as Graph Neural Networks (GNNs) to model topological relationships across complex terrain and Convolutional Long Short-Term Memory (ConvLSTM) networks to track spatiotemporal shifts in vegetation stress. Ultimately, the objective extends beyond engineering a more accurate prediction model; it is about building a scalable system that translates raw environmental data into timely, actionable intelligence—giving communities a vital advantage before a disaster strikes.

6 References

[1] NASA, “Wildfires and Climate Change - NASA Science.” Accessed: Jul. 22, 2026. [Online]. Available: https://science.nasa.gov/earth/explore/wildfires-and-climate-change/

[2] “Wildfires and Climate Change - Center for Climate and Energy SolutionsCenter for Climate and Energy Solutions.” Accessed: Jul. 22, 2026. [Online]. Available: https://www.c2es.org/content/wildfires-and-climate-change/

[3] “Wildfire Causes and Evaluations (U.S. National Park Service).” Accessed: Jul. 22, 2026. [Online]. Available: https://www.nps.gov/articles/wildfire-causes-and-evaluation.htm

[4] S. Hantson, N. Andela, M. L. Goulden, and J. T. Randerson, “Human-ignited fires result in more extreme fire behavior and ecosystem impacts,” *Nat. Commun.*, vol. 13, no. 1, p. 2717, Dec. 2022, doi: 10.1038/S41467-022-30030-2.

[5] “INCIDENT REMOTE AUTOMATIC WEATHER STATION (IRAWS) | Wildland Fire Application Information Portal.” Accessed: Jul. 22, 2026. [Online]. Available: https://www.wildfire.gov/page/incident-remote-automatic-weather-station-iraws

[6] J. Gu *et al.*, “Recent Advances in Convolutional Neural Networks”.

[7] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition”, Accessed: Jul. 22, 2026. [Online]. Available: http://image-net.org/challenges/LSVRC/2015/

[8] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature Pyramid Networks for Object Detection”.

[9] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization”, Accessed: Jul. 22, 2026. [Online]. Available: http://gradcam.cloudcv.org

[10] “Multispectral Imaging | Tree San Diego.” Accessed: Jul. 22, 2026. [Online]. Available: https://www.treesandiego.org/multispectral-imaging

[11] “Guardian Landsat on Firewatch - NASA Science.” Accessed: Jul. 23, 2026. [Online]. Available: https://science.nasa.gov/missions/landsat/guardian-landsat-on-firewatch/

[12] “Spectral Bands and Applications - NASA Science.” Accessed: Jul. 27, 2026. [Online]. Available: https://science.nasa.gov/mission/landsat/spectral-bands-and-applications/

[13] “Landsat 8 OLI and TIRS and their uses | U.S. Geological Survey.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.usgs.gov/media/images/landsat-8-oli-and-tirs-and-their-uses

[14] “NDVI vs. EVI2 vs. NIRv: Which Index Should You Trust (and When)? | ClearSKY Knowledge.” Accessed: Jul. 23, 2026. [Online]. Available: https://clearsky.vision/knowledge/ndvi-evi-what-index-to-use

[15] N. Takhtkeshha, G. Mandlburger, F. Remondino, and J. Hyyppä, “Multispectral Light Detection and Ranging Technology and Applications: A Review,” *Sensors 2024, Vol. 24, Page 1669*, vol. 24, no. 5, p. 1669, Mar. 2024, doi: 10.3390/S24051669.

[16] “The Basics of LiDAR - Light Detection and Ranging - Remote Sensing | NSF NEON | Open Data to Understand our Ecosystems.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.neonscience.org/resources/learning-hub/tutorials/lidar-basics

[17] “Multispectral Light Detection and Ranging Technology and Applications: A Review - PMC.” Accessed: Jul. 23, 2026. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC10934921/

[18] “Landsat 8 OLI and TIRS Calibration Notices | U.S. Geological Survey.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.usgs.gov/landsat-missions/landsat-8-oli-and-tirs-calibration-notices

[19] “Landsat 8 OLI and TIRS and their uses | U.S. Geological Survey.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.usgs.gov/media/images/landsat-8-oli-and-tirs-and-their-uses

[20] “Fire FAQs—What is forest fuel, and what are fuel treatments? | OSU Extension Service.” Accessed: Jul. 23, 2026. [Online]. Available: https://extension.oregonstate.edu/catalog/em-9230-fire-faqs-what-forest-fuel-what-are-fuel-treatments

[21] “Plants respond to heat differently than humans | MU Extension.” Accessed: Jul. 23, 2026. [Online]. Available: https://extension.missouri.edu/news/plants-respond-to-heat-differently-than-humans

[22] “How Ultrasonic Waves Measure Wind Speed and Direction | DFRobot Wiki.” Accessed: Jul. 23, 2026. [Online]. Available: https://wiki.dfrobot.com/tutorial/20136/

[23] “THERMISTOR BASICS – Wavelength Electronics.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.teamwavelength.com/thermistor-basics/?srsltid=AfmBOooDk5ZGjujGO6dxY6fPLrWe8qbQfB8L9LHlClqosDS34w0CU6l1

[24] “Humidity Academy Theory 6: How Capacitive Sensors Measure Humidity.” Accessed: Jul. 23, 2026. [Online]. Available: https://www.processsensing.com/en-us/blog/capacitive-sensor-technology.htm

[25] M. Kojima, “Characteristics evaluation of MEMS atmospheric pressure sensors,” *Measurement: Sensors*, vol. 38, p. 101679, May 2025, doi: 10.1016/J.MEASEN.2024.101679.

[26] “Landsat Collection 2 Level-2 Science Products | U.S. Geological Survey.” Accessed: Jul. 24, 2026. [Online]. Available: https://www.usgs.gov/landsat-missions/landsat-collection-2-level-2-science-products

[27] Z. Jiang, A. R. Huete, K. Didan, and T. Miura, “Development of a two-band enhanced vegetation index without a blue band,” *Remote Sens. Environ.*, vol. 112, no. 10, pp. 3833–3845, Oct. 2008, doi: 10.1016/J.RSE.2008.06.006.

[28] “EVI2 (2-band Enhanced Vegetation Index) | Glossary.” Accessed: Jul. 24, 2026. [Online]. Available: https://giscarta.com/gis-glossary/evi2

[29] “Radiometric Calibration - an overview | ScienceDirect Topics.” Accessed: Jul. 24, 2026. [Online]. Available: https://www.sciencedirect.com/topics/engineering/radiometric-calibration

[30] M. T. Macdonald, T. J. Esau, and M. F. Bilodeau, “Exploring the potential of two vegetative indices, normalized difference vegetation index (NDVI) and enhanced vegetation index 2 (EVI2), to estimate balsam fir nutrition using uncrewed aerial vehicles and multispectral sensing,” *https://doi.org/10.1139/cjps-2024-0072*, vol. 105, 2025, doi: 10.1139/CJPS-2024-0072.

[31] D. G. Lee and D. C. Lee, “Improvement scheme of airborne lidar strip adjustment,” *Journal of the Korean Society of Surveying, Geodesy, Photogrammetry and Cartography*, vol. 36, no. 5, pp. 355–369, 2018, doi: 10.7848/KSGPC.2018.36.5.355.

[32] L. Ma, Z. Song, and G. Zhu, “An improved SURF algorithm based local image symmetry scoring scheme,” *Proceedings - 2014 7th International Congress on Image and Signal Processing, CISP 2014*, pp. 264–268, Jan. 2014, doi: 10.1109/CISP.2014.7003789.

[33] C. Liu, H. Liu, Y. Liu, T. Li, and T. Wang, “Normalized cross correlation image stitching algorithm based on minimum spanning tree,” *Optik (Stuttg).*, vol. 179, pp. 610–616, Feb. 2019, doi: 10.1016/J.IJLEO.2018.10.166.

[34] M. A. Fischler and R. C. Bolles, “Random sample consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography,” *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981, doi: 10.1145/358669.358692;PAGE:STRING:ARTICLE/CHAPTER.

[35] "Complete PPK Workflow for DJI Enterprise Drones." Accessed: Jul. 24, 2026. [Online]. Available: https://enterprise-insights.dji.com/blog/ppk-post-processed-kinematics-workflow

[36] "LiDAR Boresight Calibration and Strip Alignment – Part 2 - AEVEX Geodetics." Accessed: Jul. 24, 2026. [Online]. Available: https://geodetics.com/strip-alignment/

[37] K. Liu, H. Ma, L. Zhang, Z. Cai, and H. Ma, "Strip Adjustment of Airborne LiDAR Data in Urban Scenes Using Planar Features by the Minimum Hausdorff Distance," *Sensors 2019, Vol. 19, Page 5131*, vol. 19, no. 23, p. 5131, Nov. 2019, doi: 10.3390/S19235131.

[38] "Removing outliers using a StatisticalOutlierRemoval filter — Point Cloud Library 1.15.1-dev documentation." Accessed: Jul. 24, 2026. [Online]. Available: https://pointclouds.org/documentation/tutorials/statistical_outlier.html

[39] S. Pan *et al.*, "Land-cover classification of multispectral LiDAR data using CNN with optimized hyper-parameters," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 166, pp. 241–254, Aug. 2020, doi: 10.1016/J.ISPRSJPRS.2020.05.022.

[40] B. J. Pickstone, H. A. Graham, and A. M. Cunliffe, "Estimating canopy height in tropical forests: Integrating airborne LiDAR and multi-spectral optical data with machine learning," *Sustainable Environment*, vol. 11, no. 1, Dec. 2025, doi: 10.1080/27658511.2025.2469406.

[41] E. Wang, H. Wan, A. Simeoni, C. Mou, and Y. Zhang, "Experimental Study on Effects of Moisture Content and Tree Height on Crown Fire Behaviors of Live Cypress Trees," *Combustion Science and Technology*, vol. 197, no. 11, pp. 2727–2742, Aug. 2025, doi: 10.1080/00102202.2024.2323576.

[42] "Voxelization - Open3D 0.19.0 documentation." Accessed: Jul. 24, 2026. [Online]. Available: https://www.open3d.org/docs/release/tutorial/geometry/voxelization.html

[43] J. Yang, L. El Mendili, Y. Khayer, S. McArdle, and L. H. Beni, "INSTANCE SEGMENTATION OF LIDAR DATA WITH VISION TRANSFORMER MODEL IN SUPPORT INUNDATION MAPPING UNDER FOREST CANOPY ENVIRONMENT," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLVIII-1-W2-2023, no. 1/W2-2023, pp. 203–208, Dec. 2023, doi: 10.5194/ISPRS-ARCHIVES-XLVIII-1-W2-2023-203-2023.

[44] "Fuels: Canopy Fuel Characteristics | NWCG." Accessed: Jul. 24, 2026. [Online]. Available: https://www.nwcg.gov/publications/pms437/fuels/canopy-fuel-characteristics

[45] R. Wang *et al.*, "Quantifying Understory Vegetation Cover of Pinus massoniana Forest in Hilly Region of South China by Combined Near-Ground Active and Passive Remote Sensing," *Drones 2022, Vol. 6, Page 240*, vol. 6, no. 9, p. 240, Sep. 2022, doi: 10.3390/DRONES6090240.

[46] "What are Ladder Fuels? | Orinda, CA." Accessed: Jul. 24, 2026. [Online]. Available: https://www.cityoforinda.gov/753/What-are-Ladder-Fuels

[47] “Min-Max Normalization. Instead of directly jumping into the… | by Abhinav Bandaru | Medium.” Accessed: Jul. 24, 2026. [Online]. Available: https://medium.com/@abhi1achiever/min-max-normalization-db1f515b08b4

[48] C. Huang *et al.*, “Automated masking of cloud and cloud shadow for forest change analysis using Landsat images,” *Int. J. Remote Sens.*, vol. 31, no. 20, pp. 5449–5464, 2010, doi: 10.1080/01431160903369642.

[49] “Multilayer Perceptrons in Machine Learning: A Comprehensive Guide | DataCamp.” Accessed: Jul. 24, 2026. [Online]. Available: https://www.datacamp.com/tutorial/multilayer-perceptrons-in-machine-learning

[50] “Displaying Speed and Direction Symbology from U and V vectors.” Accessed: Jul. 24, 2026. [Online]. Available: https://www.esri.com/arcgis-blog/products/product/analytics/displaying-speed-and-direction-symbology-from-u-and-v-vectors

[51] “Computing wind direction and speed from u and v.” Accessed: Jul. 25, 2026. [Online]. Available: http://www.weatherclasses.com/uploads/1/3/1/3/131359169/computing_wind_direction_and_speed_from_u_and_v.pdf

[52] L. Li, B. Chen, D. Duan, and L. Liu, “High-Precision Time Synchronization Based on Timestamp Mapping in Datacenter Networks,” *Electronics 2025, Vol. 14, Page 610*, vol. 14, no. 3, p. 610, Feb. 2025, doi: 10.3390/ELECTRONICS14030610.

[53] “What is Forward and Backward fillling of Imputation in Exploratory Data Analysis (EDA)? | by Sujatha Mudadla | Medium.” Accessed: Jul. 25, 2026. [Online]. Available: https://medium.com/@sujathamudadla1213/what-is-forward-and-backward-fillling-of-imputation-in-exploratory-data-analysis-eda-74ec33aca4bd

[54] “Convolutional Neural Network (CNN) in Deep Learning - GeeksforGeeks.” Accessed: Jul. 25, 2026. [Online]. Available: https://www.geeksforgeeks.org/deep-learning/convolutional-neural-network-cnn-in-machine-learning/#google_vignette

[55] “What is Backpropagation? | IBM.” Accessed: Jul. 25, 2026. [Online]. Available: https://www.ibm.com/think/topics/backpropagation

[56] K. Fedorchuk, S. M. Russell, K. Zibaei, M. Yassin, and D. G. Hicks, “DeepKymoTracker: A tool for accurate construction of cell lineage trees for highly motile cells,” *bioRxiv*, p. 2024.02.28.582654, Mar. 2024, doi: 10.1101/2024.02.28.582654.

[57] “Measuring Vegetation (NDVI & EVI) - NASA Science.” Accessed: Jul. 25, 2026. [Online]. Available: https://science.nasa.gov/earth/earth-observatory/measuring-vegetation-ndvi-evi/

[58] A. Chroni, C. Vasilakos, M. Christaki, and N. Soulakellis, “Fusing Multispectral and LiDAR Data for CNN-Based Semantic Segmentation in Semi-Arid Mediterranean Environments: Land Cover Classification and Analysis,” *Remote Sensing 2024, Vol. 16, Page 2729*, vol. 16, no. 15, p. 2729, Jul. 2024, doi: 10.3390/RS16152729.

[59] L. Cao, S. Gao, P. Li, T. Yun, X. Shen, and H. Ruan, "Aboveground Biomass Estimation of Individual Trees in a Coastal Planted Forest Using Full-Waveform Airborne Laser Scanning Data," *Remote Sensing 2016, Vol. 8, Page 729*, vol. 8, no. 9, p. 729, Sep. 2016, doi: 10.3390/RS8090729.

[60] L. Giglio, W. Schroeder, and C. O. Justice, "The collection 6 MODIS active fire detection algorithm and fire products-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)," *Remote Sens. Environ.*, vol. 178, pp. 31–41, 2016, doi: 10.1016/j.rse.2016.02.054.

[61] M. D. Zeiler and R. Fergus, "Visualizing and Understanding Convolutional Networks arXiv:1311.2901v3 [cs.CV] 28 Nov 2013," *Computer Vision–ECCV 2014*, vol. 8689, no. PART 1, pp. 818–833, 2014, doi: 10.1007/978-3-319-10590-1_53.

[62] S. P. H. Boroujeni *et al.*, "A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management," *Information Fusion*, vol. 108, p. 102369, Aug. 2024, doi: 10.1016/J.INFFUS.2024.102369.

[63] K. C. Short, "Spatial wildfire occurrence data for the United States, 1992-2020 [FPA_FOD_20221014]," *Forest Service Research Data Archive*, doi: 10.2737/RDS-2013-0009.6.

[64] "Climate Data Online (CDO) - The National Climatic Data Center's (NCDC) Climate Data Online (CDO) provides free access to NCDC's archive of historical weather and climate data in addition to station history information. | National Climatic Data Center (NC…." Accessed: Jul. 26, 2026. [Online]. Available: https://www.ncei.noaa.gov/cdo-web/

[65] T. Alhajj and V. Corlay, "Improved Contact Graph Routing in Delay Tolerant Networks with Capacity and Buffer Constraints," Oct. 2024, Accessed: Jul. 26, 2026. [Online]. Available: https://arxiv.org/pdf/2410.15546

[66] "USA: How drones and AI are changing the way wildfires are fought | PreventionWeb." Accessed: Jul. 26, 2026. [Online]. Available: https://www.preventionweb.net/news/how-drones-and-ai-are-changing-way-wildfires-are-fought

[67] S. P. H. Boroujeni *et al.*, "A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management," *Information Fusion*, vol. 108, p. 102369, Aug. 2024, doi: 10.1016/J.INFFUS.2024.102369.

[68] "AI-Enabled Drone Swarms for Fire Detection, Mapping, and Modeling - NASA Earth Science and Technology Office." Accessed: Jul. 26, 2026. [Online]. Available: https://esto.nasa.gov/firetech/ai-enabled-drone-swarms-for-fire-detection-mapping-and-modeling/

[69] Y. Li, T. Zhang, Y. Ding, R. Wadhwani, and X. Huang, "Review and perspectives of digital twin systems for wildland fire management," *Journal of Forestry Research 2024 36:1*, vol. 36, no. 1, pp. 14-, Dec. 2024, doi: 10.1007/S11676-024-01810-X.

[70] "NASA 'Wildfire Digital Twin' Pioneers New AI Models and Streaming Data Techniques for Forecasting Fire and Smoke - NASA Science." Accessed: Jul. 26, 2026. [Online]. Available: https://science.nasa.gov/science-research/science-enabling-technology/nasa-

wildfire-digital-twin-pioneers-new-ai-models-and-streaming-data-techniques-for-forecasting-fire-and-smoke/

[71] K. Vogiatzoglou, C. Papadimitriou, V. Bontozoglou, and K. Ampountolas, “Physics-informed neural networks for parameter learning of wildfire spreading,” *Comput. Methods Appl. Mech. Eng.*, vol. 434, p. 117545, Feb. 2025, doi: 10.1016/J.CMA.2024.117545.

[72] “(PDF) A Unified Spatio-Temporal Deep Learning Framework for Probabilistic Wildfire Ignition Forecasting Using Satellite Data Fusion and 3D Canopy Metrics.” Accessed: Jul. 26, 2026. [Online]. Available: https://www.researchgate.net/publication/399422303_A_Unified_Spatio-Temporal_Deep_Learning_Framework_for_Probabilistic_Wildfire_Ignition_Forecasting_Using_Satellite_Data_Fusion_and_3D_Canopy_Metrics

[73] J. J. Dabrowski *et al.*, “Bayesian Physics Informed Neural Networks for data assimilation and spatio-temporal modelling of wildfires,” *Spat. Stat.*, vol. 55, p. 100746, Jun. 2023, doi: 10.1016/J.SPASTA.2023.100746.

[74] B. Shaddy *et al.*, “Generative Algorithms for Fusion of Physics-Based Wildfire Spread Models with Satellite Data for Initializing Wildfire Forecasts,” *Artificial Intelligence for the Earth Systems*, vol. 3, no. 3, Jul. 2024, doi: 10.1175/AIES-D-23-0087.1.

[75] F. Huot *et al.*, “Deep Learning Models for Predicting Wildfires from Historical Remote-Sensing Data,” Oct. 2020, Accessed: Jul. 26, 2026. [Online]. Available: https://arxiv.org/pdf/2010.07445