VI-NET: A HYBRID DEEP LEARNING MODEL FOR WILDFIRE SPREAD PREDICTION AND OPTIMIZED SAFE PATH PLANNING

Manavjit Singh Dhindsa and **Kshirasagar (Sagar) Naik,** Department of ECE

Some results from an NSERC (Natural Sciences and Engineering Research Council of Canada) project on forest fire modeling.

Universities involved: Waterloo, Carleton, Dalhousie

DBKDA 2025, Lisbon, March 9-13, 2025

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

IARIA

Agenda:

- Introduction
- Background and Literature Review
- Problem Formulation
- Dataset Description and Preparation
- Methodology
- Results and Experiments
- Conclusion and Future Work

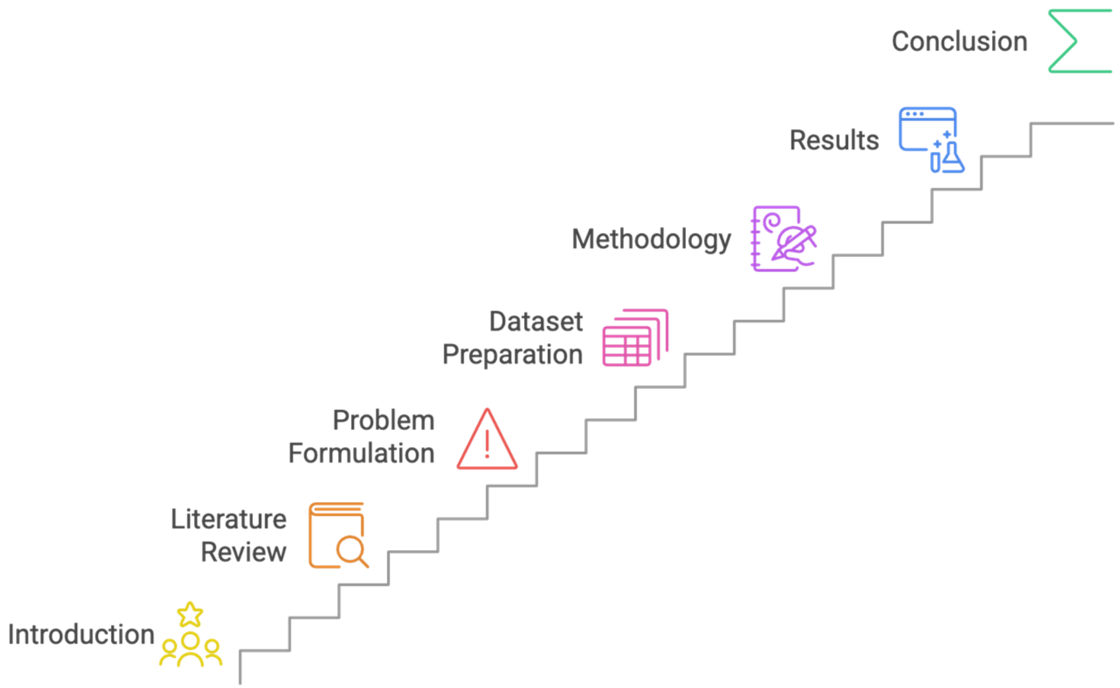

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

INTRODUCTION

Introduction: Motivation

Motivation

- Wildfires are increasing in frequency and intensity globally.
- Climate change and human activities are major contributors.

  - human activity and lightning: almost 50:50

- 2020: Over 4M hectares burned in the U.S., $19B economic loss.
- Releases carbon and pollutants, worsening air quality.
- Threatens biodiversity, infrastructure, and human lives.

Introduction : LA Fires 2025

- Devastating wildfires raged across LA since Jan 7, 2025, fueled by severe Santa Ana winds and dry conditions. [1]
- At least 18 deaths reported, with over 180,000 people evacuated. More than 13,400 structures destroyed or damaged, scorched over 200,000 acres of land. [1]
- Major challenge: Fires encroaching narrow, winding roads in affluent suburbs hindering quick evacuations, causing gridlock [2]
- Critical Issue: Inefficient resource allocation led to dry hydrants and low water pressure in several areas limiting firefighters' ability to combat the blazes effectively. [3]

[1] https://jnylaw.com/blog/firefighting-efforts-in-the-2025-los-angeles-wildfire/
[2] https://www.bbc.com/future/article/20250109-why-los-angeles-was-so-hard-to-evacuate-during-the-wildfires/
[3] https://www.latimes.com/environment/story/2025-01-09/california-fires-water-supply-problems

Introduction: Problem Statement

Wildfires are dynamic and complex natural disasters that pose a significant challenge for prediction and management. Traditional models often struggle to capture the spatial and temporal dependencies critical for accurate fire spread prediction. Enhanced prediction models are essential due to the increasing severity and frequency of wildfires.

The aim is to predict the next day's wildfire spread using a set of environmental inputs, such as weather and topographical data, modeled as a semantic segmentation problem.

The task also involves utilizing this prediction to compute safe paths that safely avoid the fire-affected zones, thereby aiding in effective disaster management and evacuation strategies.

INTRODUCTION: OBJECTIVES

Develop Vi-Net: Introduce a novel hybrid deep learning model combining the **U-Net** and **Vision Transformer (ViT)** architectures to leverage their complementary strengths.

Predict Wildfire Spread: Achieve high accuracy and recall in **next-day** wildfire spread predictions using a multimodal dataset.

Integrate Safe Path Planning: Incorporate wildfire predictions into the **A*** algorithm for generating optimized, safe evacuation routes in fire-prone areas, addressing the urgent need for adaptive disaster management solutions.

Address Data Imbalance: Utilize advanced loss functions like **Focal Tversky Loss** to prioritize minority classes (fire regions) and improve prediction sensitivity.

Enable Generalization: Ensure the model's applicability across diverse geographical regions and unseen datasets for robust real-world implementation.

Lay the Foundation for Future Work: Set the groundwork for integrating real-time data and expanding to global wildfire datasets.

Introduction: System Model

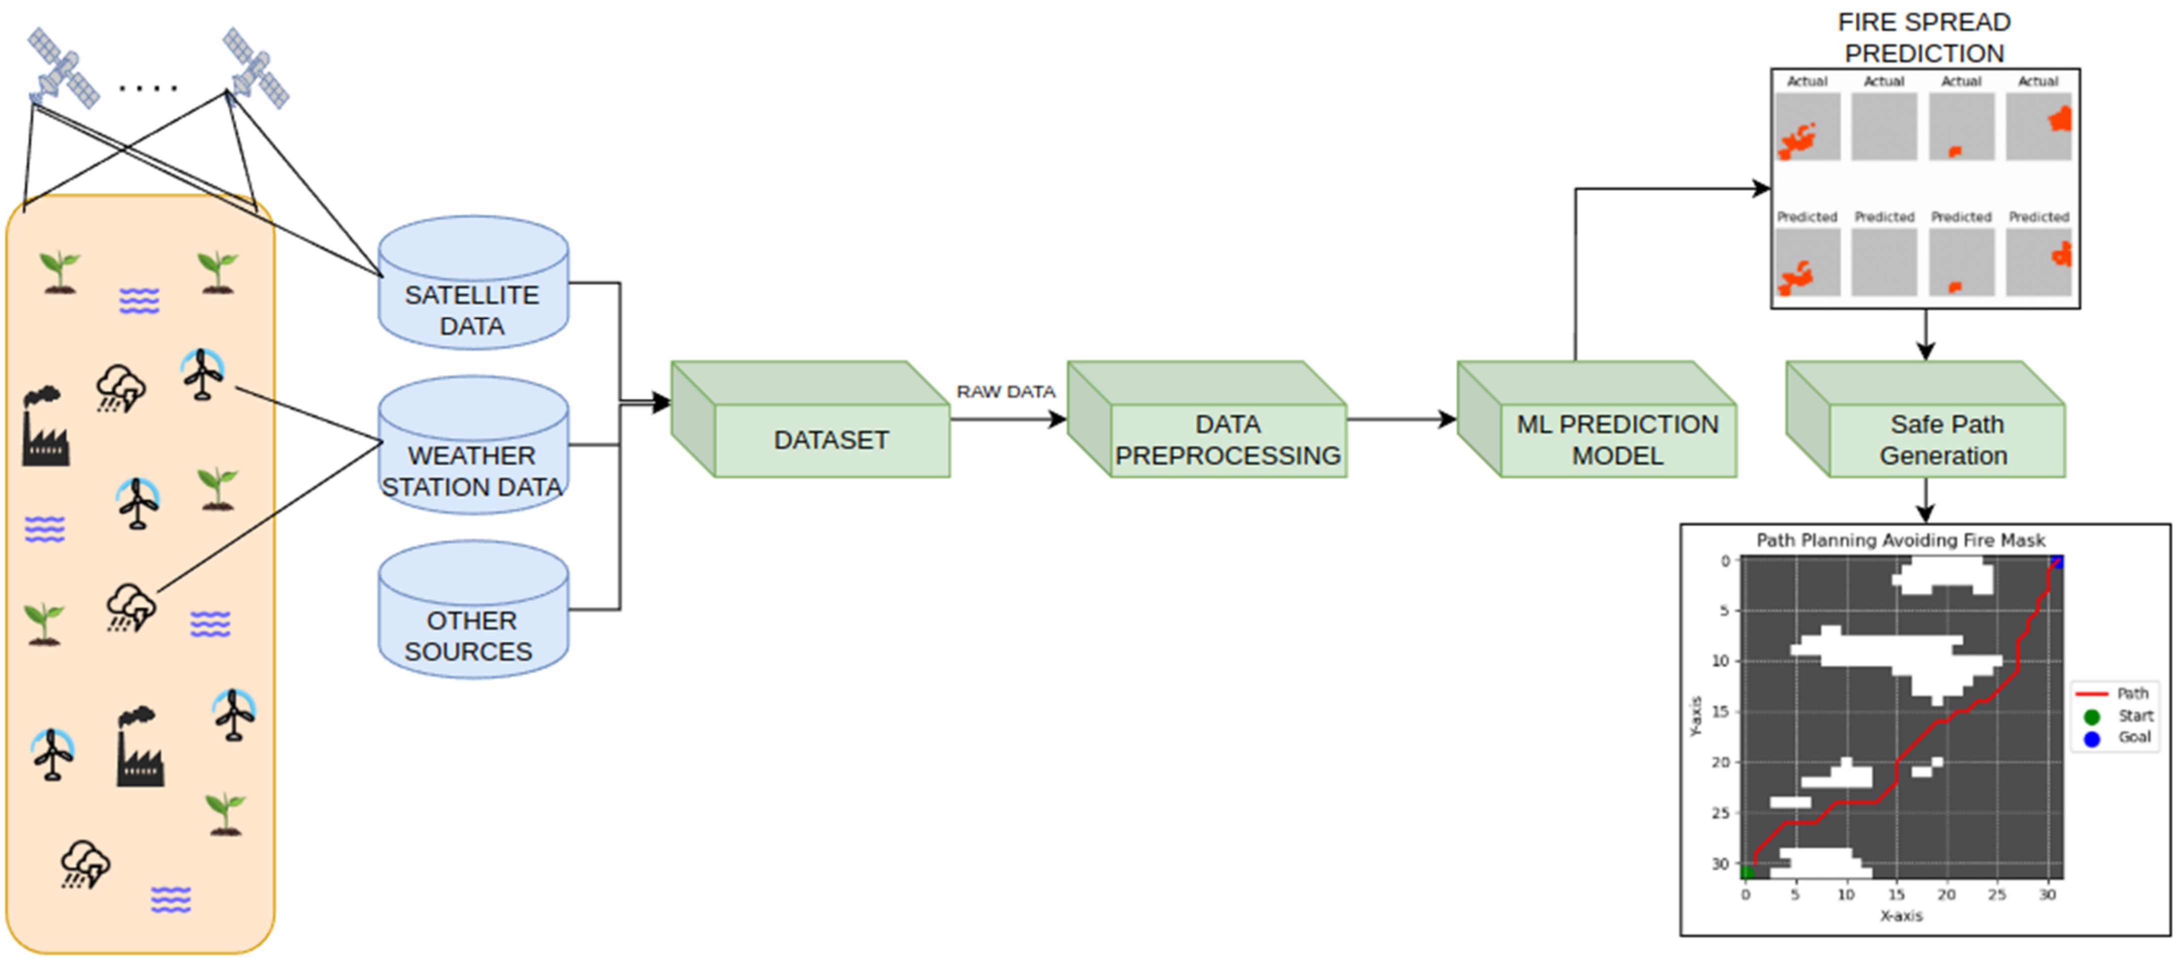

Background and Literature Review

Background

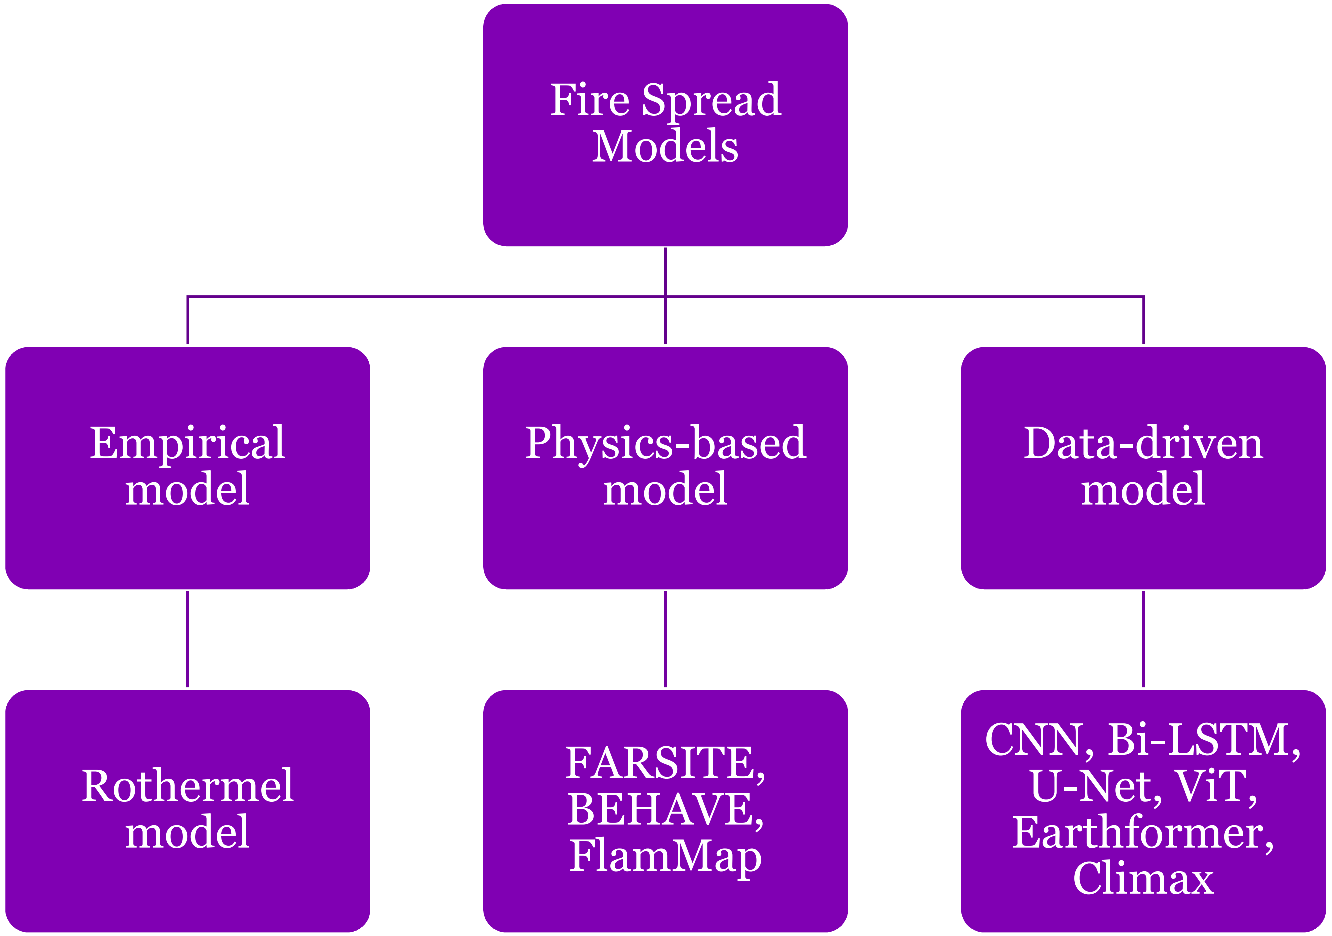
Fire Spread Models
Empirical model
Physics-based model
Data-driven model
Rothermel model
FARSITE, BEHAVE, FlamMap
CNN, Bi-LSTM, U-Net, ViT, Earthformer, Climax

Background

Ref No	Paper Title	Year	Journal	Data Sources	Approach
[3]	The **Rothermel** Surface Fire Spread Model and Associated Developments	2018	Comprehensive Explanation Report	Historical wildfire events	Empirical modeling
[17]	**FARSITE:** Fire Area Simulator-model development and evaluation	1998	US Forest Service	Topographical and meteorological data	Physics-based modeling with spatio-temporal predictions
[18]	An Overview of **FlamMap** Fire Modeling Capabilities	2006	Conference Proceedings	Geospatial and meteorological data	Integrated semi-empirical fire modeling
[33]	A Model for Predicting Forest Fire Spreading Using **Cellular Automata**	1997	Ecological Modelling	Cellular automata grids	Stochastic simulations using cellular automata
[22]	Wildfire Segmentation Using Deep **Vision Transformers**	2021	Remote Sensing	Satellite imagery and meteorological data	Deep vision transformers for segmentation
[1]	Emulation of Wildland Fire Spread Simulation Using **Deep Learning**	2021	Neural Networks	Deep learning models and wildfire simulators	Deep learning emulation of fire spread

Background

[28]	Hexagonal Cellular Automaton Model for Fire Spread Simulation	2007	Fire Safety Journal	Hexagonal grid-based fire data	Enhanced spatial modeling via hexagonal automata
[12]	Using SVM for Forest Fire Prediction	2019	Journal of Environmental Management	Environmental metrics and SVM outputs	Support Vector Machines for prediction
[32]	A Review of Machine Learning Applications in Wildfire Science and Management	2020	CoRR	Various wildfire-related datasets	Comprehensive ML review and applications
[19]	Paying Attention to Wildfire: Using U-Net with Attention Blocks on Multimodal Data	2023	ICMI Proceedings	Satellite and ground-based sensors	Attention-enhanced U-Net for prediction
[23]	Physics-based Model of Wildfire Propagation Towards Faster-than-Real-Time Simulations	2020	Computers & Mathematics with Applications	Physics-based simulations	Physics-based faster-than-real-time modeling

Background

[21]	WildfireSpreadTS: A Dataset of Multimodal Time Series for Wildfire Spread Prediction	2023	Advances in Neural Information Processing Systems	Multimodal time-series data	Time-series analysis for spread prediction
[4]	Machine Learning and Deep Learning for Wildfire Spread Prediction: A Review	2024	Fire	Historical fire data and environmental models	Comprehensive DL review in wildfire spread
[20]	Earthformer: Exploring Space-Time Transformers for Earth System Forecasting	2022	Advances in Neural Information Processing Systems	Global earth system data	Space-time transformers for forecasting
[11]	Extreme Fire Spread Events and Area Burned under Recent and Future Climate in the Western USA	2022	Global Ecology and Biogeography	Climate models and ecological metrics	Ecological and biogeographical insights on fire
[39]	CNN-BiLSTM: A novel deep learning model for near-real-time daily wildfire spread prediction.	2024	Remote Sensing	Geospatial weather data	CNN + LSTM networks
[38]	Firepred: A hybrid multi-temporal CNN model for wildfire spread prediction	2023	Ecological Informatics	Satellite and ground-based sensors	Multi-temporal convolutional neural network

Problem Formulation

Problem Formulation

1. **Critical Need for Wildfire Prediction:** Wildfire prediction safeguards lives, property, and ecosystems by addressing the dynamic interplay of environmental, meteorological, and human factors.
2. **Leveraging Time-Series and Causality:** Time-series models capture temporal dependencies (e.g., today's weather shaping tomorrow's fire) and spatial correlations (e.g., wind affecting nearby areas), emphasizing causative relationships.
3. **One-Day-Ahead Forecasting Approach:** Focused on predicting wildfire spread for the next day, this method combines past day's observations to align with time-series principles and ensure actionable insights.
4. **Challenges in Modeling Wildfire Dynamics:** Accurate forecasting requires handling high-dimensional, multivariate, multimodal, and temporal data to capture interactions between variables like wind speed, humidity, and vegetation dryness.

Problem Formulation

- **Core Problems:**

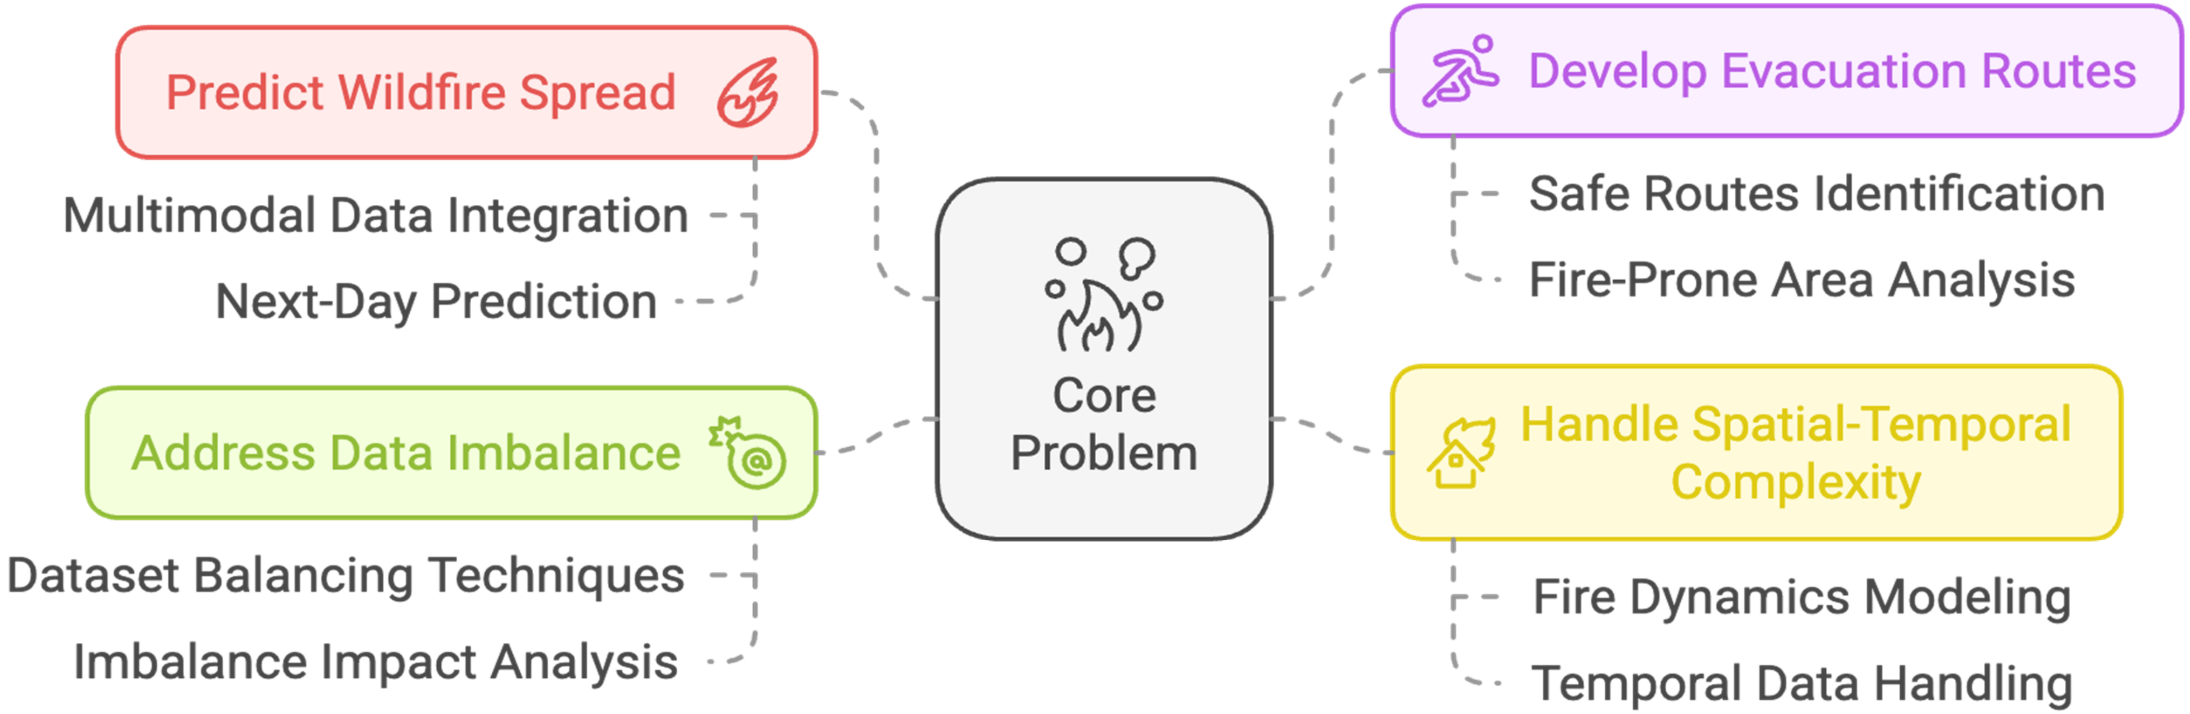

Problem Formulation: Mathematical Formulation (Wildfire Spread)

- *<u>Input:</u> Multimodal tensor representing spatial-temporal features.*
- **Framing the Problem**:
  - Predict next-day wildfire spread as a semantic segmentation task.
- **Input Representation**:
  - input structured as a 3D tensor $X \in R^{H \times W \times C}$:
    - H,W: Spatial dimensions (region height and width).
    - C: Feature channels (e.g., wind speed, temperature, vegetation indices).
  - $x(h,w,f) \in X$ represents the value of feature f at spatial location (h, w).
- **Output Objective**:
  - Predict a binary segmentation mask $Y \in R^{H \times W}$:
    - y(h,w) >= threshold: Fire presence.
    - y(h,w) < threshold: Fire absence.
- **Mapping Function**:
  - A deep learning model F(X) maps input X to output Y:
    - $F : X \rightarrow Y$
  - Captures complex interdependencies and temporal patterns in wildfire dynamics.
- *<u>Output:</u> Binary fire mask predicting next-day fire spread.*

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Problem Formulation: Model Scope

- **U-Net**: Extracts localized spatial details for precise boundary detection.
- **ViT**: Provides global context for broader fire spread understanding.
- **Vi-Net**: Combines both for robust, fine-grained accuracy and contextual coherence.

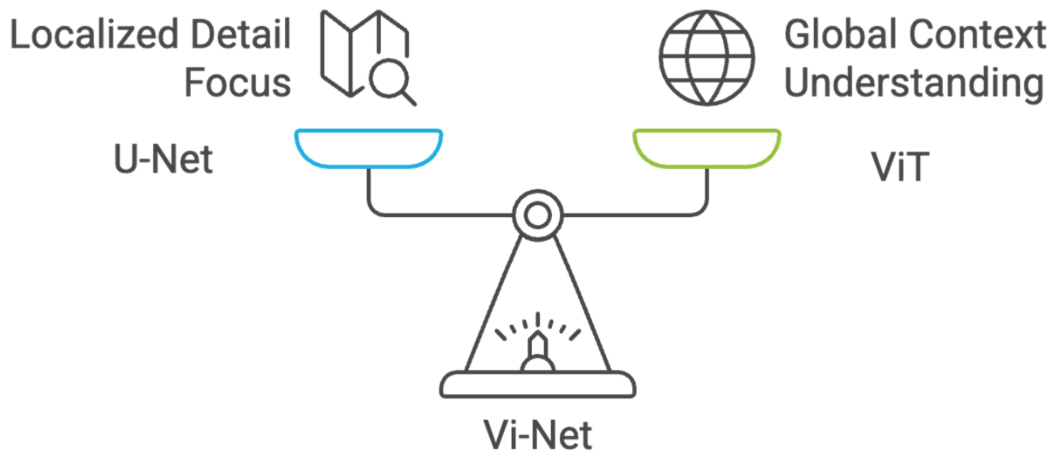

Problem Formulation : Dataset Sample

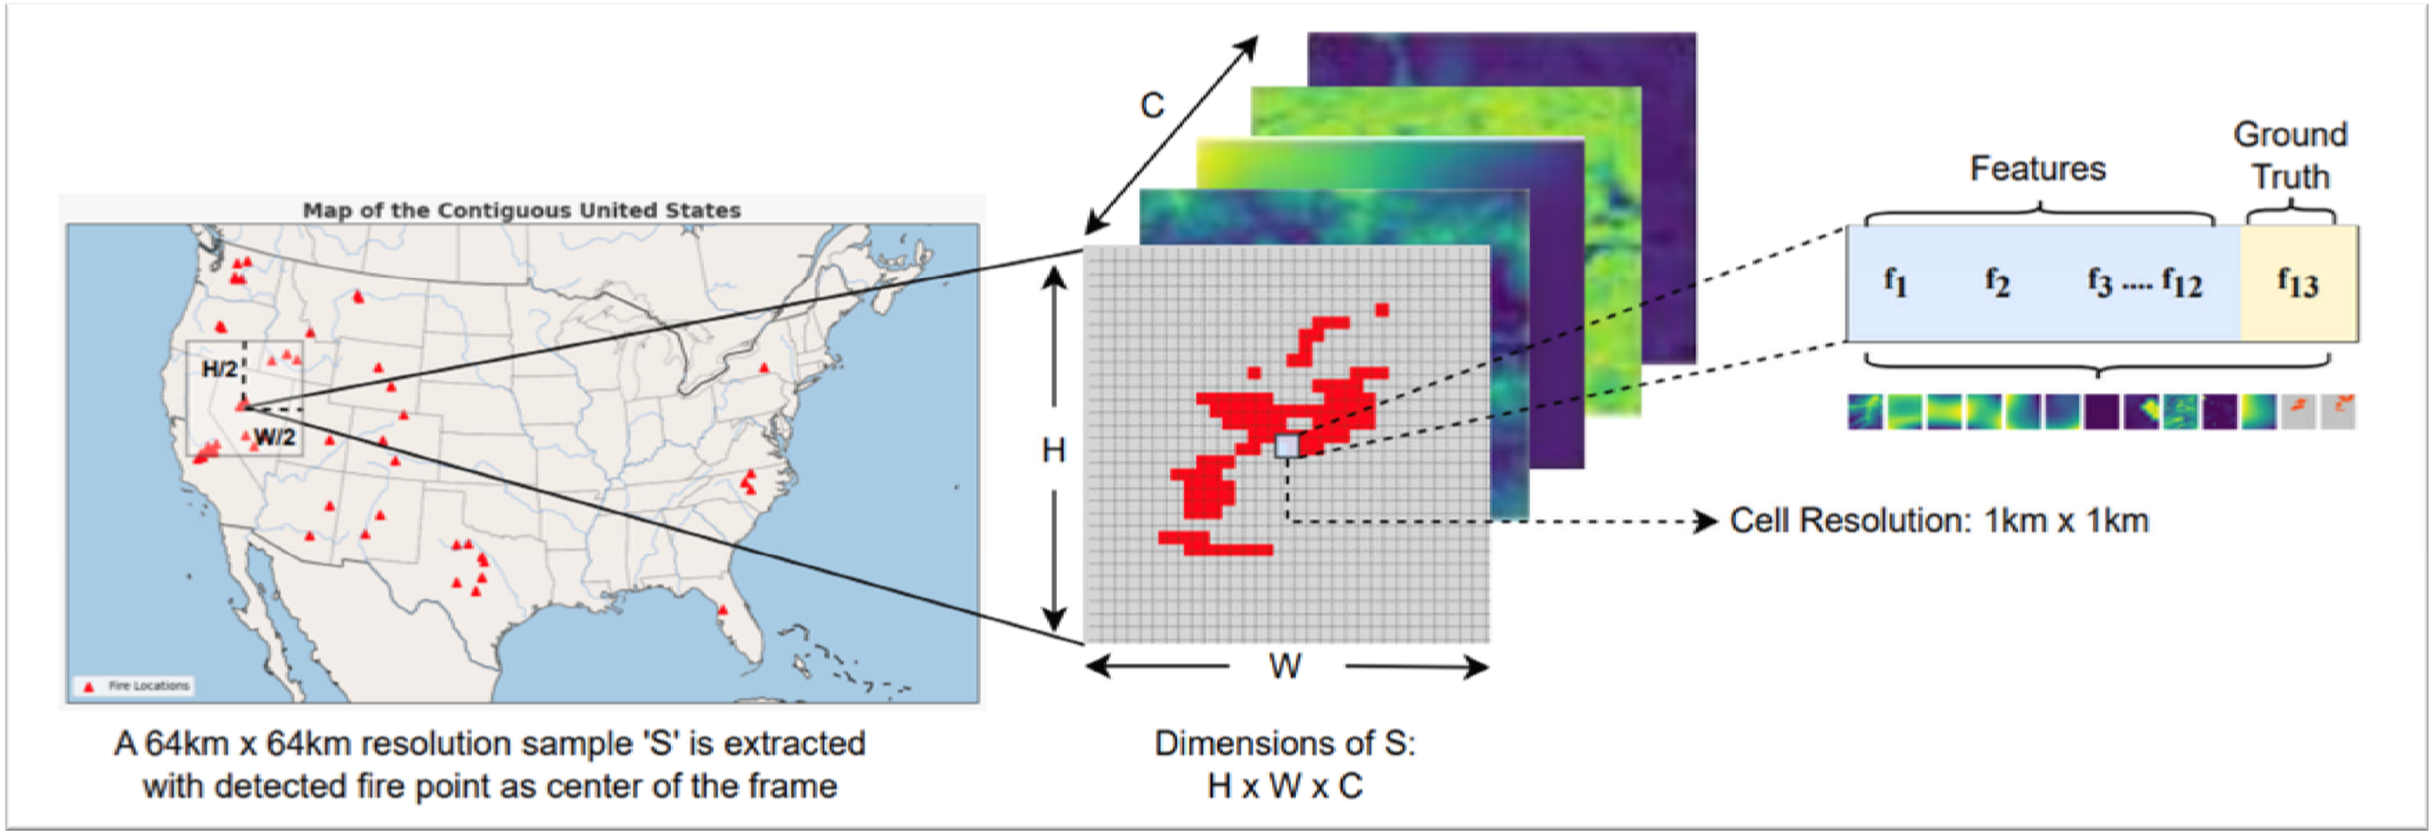

Grid Dimensions: 64 km × 64 km
Resolution: 1 km × 1 km

Problem Formulation: Path Planning Task Formulation

Output from wildfire spread prediction is used as input for safe path planning. Ensures predictions directly guide evacuation strategies.

Binary Fire Masks:

- Fire predictions are converted into binary grids.
- Each cell in the grid is classified as:
  - **1**: Fire-prone (non-passable).
  - **0**: Safe (passable).

Optimized A* Algorithm (OA*):

- Applies the binary fire mask as a grid for pathfinding.
- Computes the shortest and safest route from source to destination.
- Avoids high-risk regions identified in the fire mask.

Problem Formulation: Model Scope

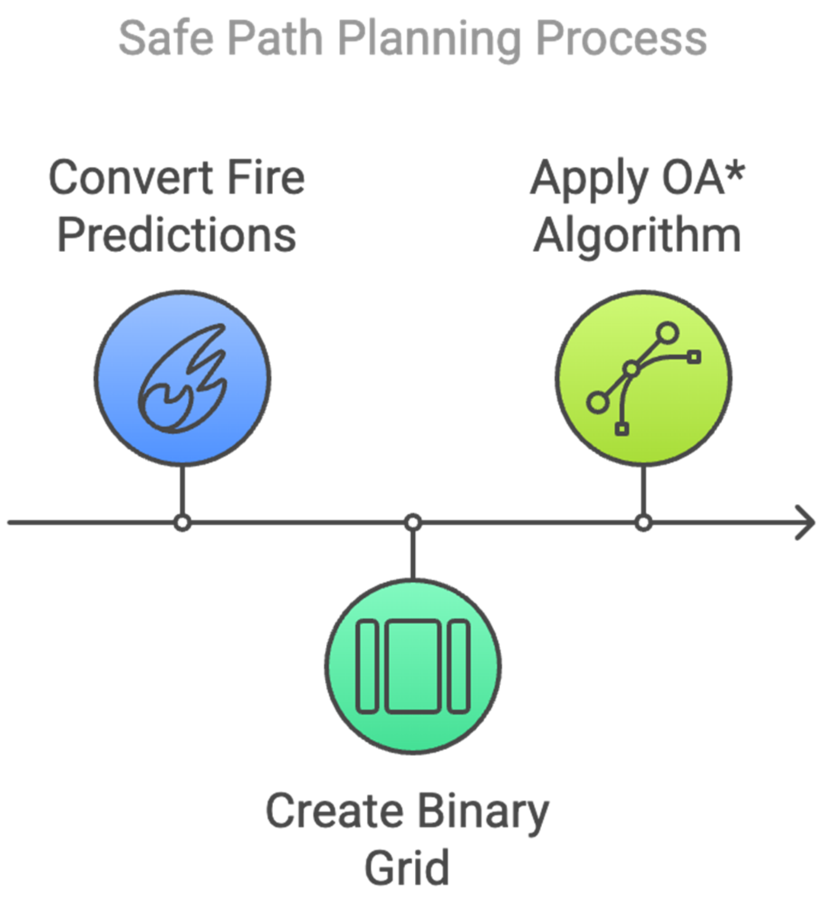

- Fire predictions are transformed into a binary mask.
- A grid is created to represent the fire mask.
- The optimized A* algorithm is applied to the grid.

Problem Formulation: Binary Grid

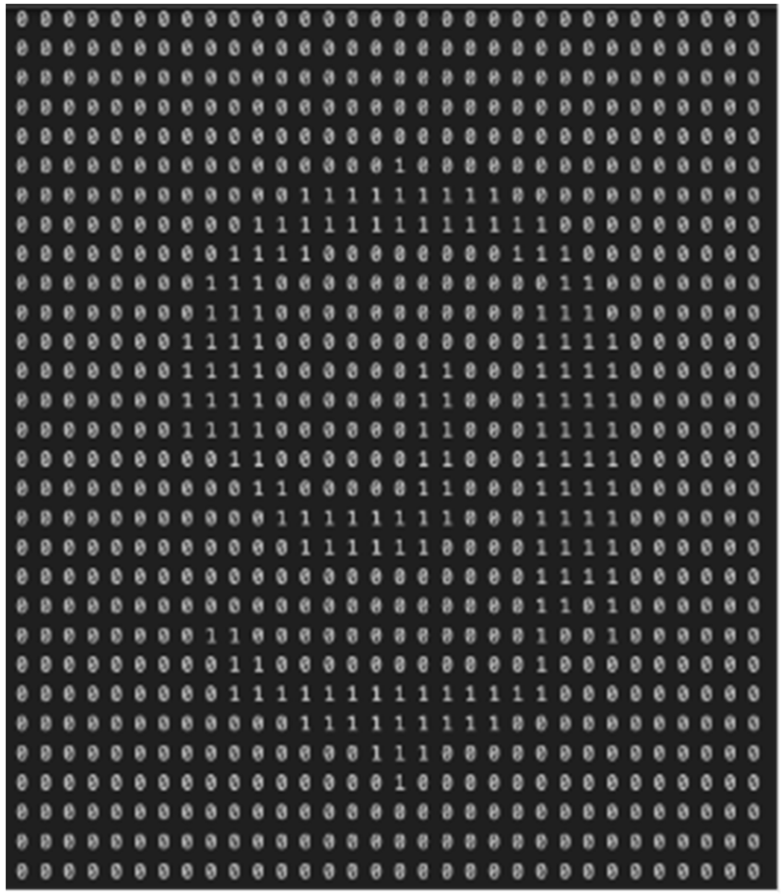

Binary Grid

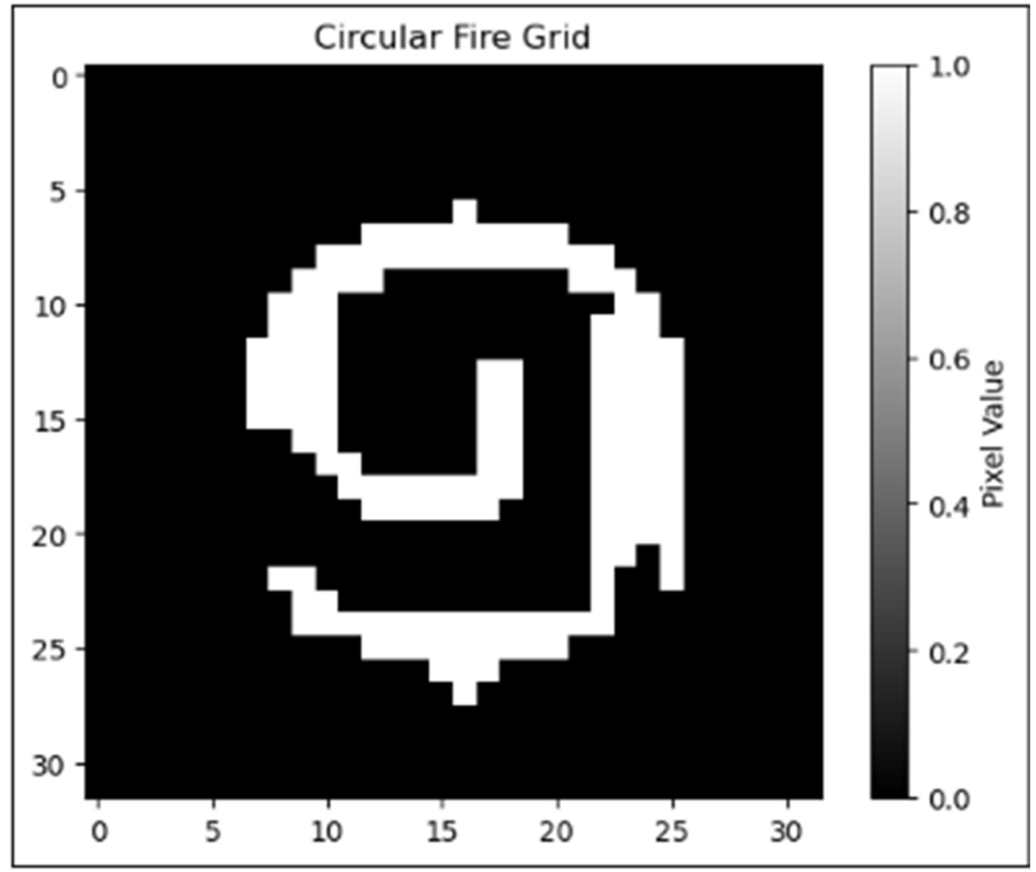

Corresponding Fire Mask

Dataset Description and Preparation

Dataset: Key Points

Dataset: Next Day Wildfire Spread

- Covers U.S. wildfire data from 2012 to 2020.
- Aggregated using the Google Earth Engine (GEE) framework, enabling the capture of a large ensemble of fire events and observational variable.
- 18,545 recorded fire events for robust analysis.
- **Multivariate Data**: Contains 13 features, consists of 11 observational variables,
- Provides two snapshots of fire spread: at time t and t+1 day, allowing for the analysis of fire dynamics over time.
- 1 km spatial resolution for detailed insights.
- Dataset split: 80% training, 10% validation, 10% testing.

Dataset: Data Collection Framework

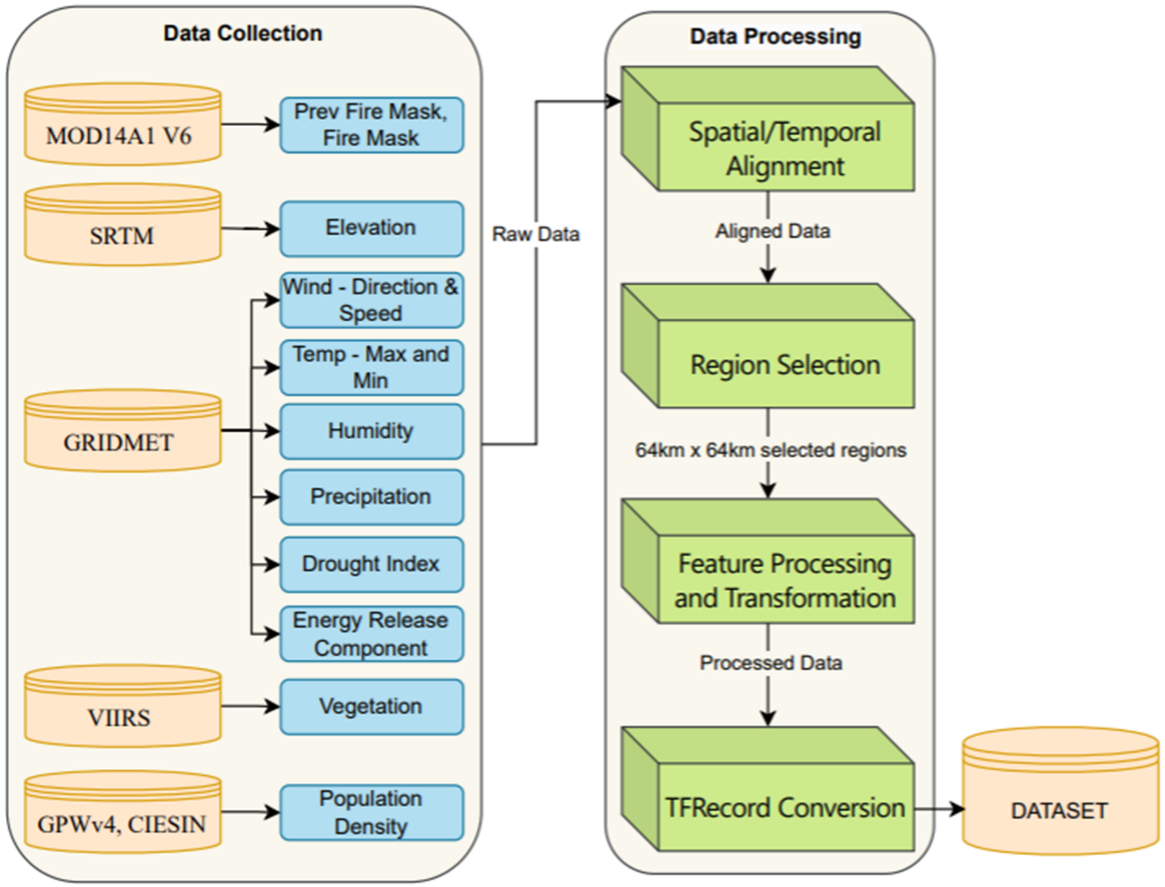

Multimodal Dataset

- **MODIS**: Fire masks at 1 km resolution.
- **GRIDMET**: Weather (temperature, wind, precipitation, humidity) and drought index at 4 km resolution.
- **VIIRS**: Vegetation indices at 0.5 km resolution.
- **SRTM**: Elevation data at 30 m resolution.
- **GPWv4:** Delivered population density at 1 km resolution.

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Dataset: Data Collection Framework

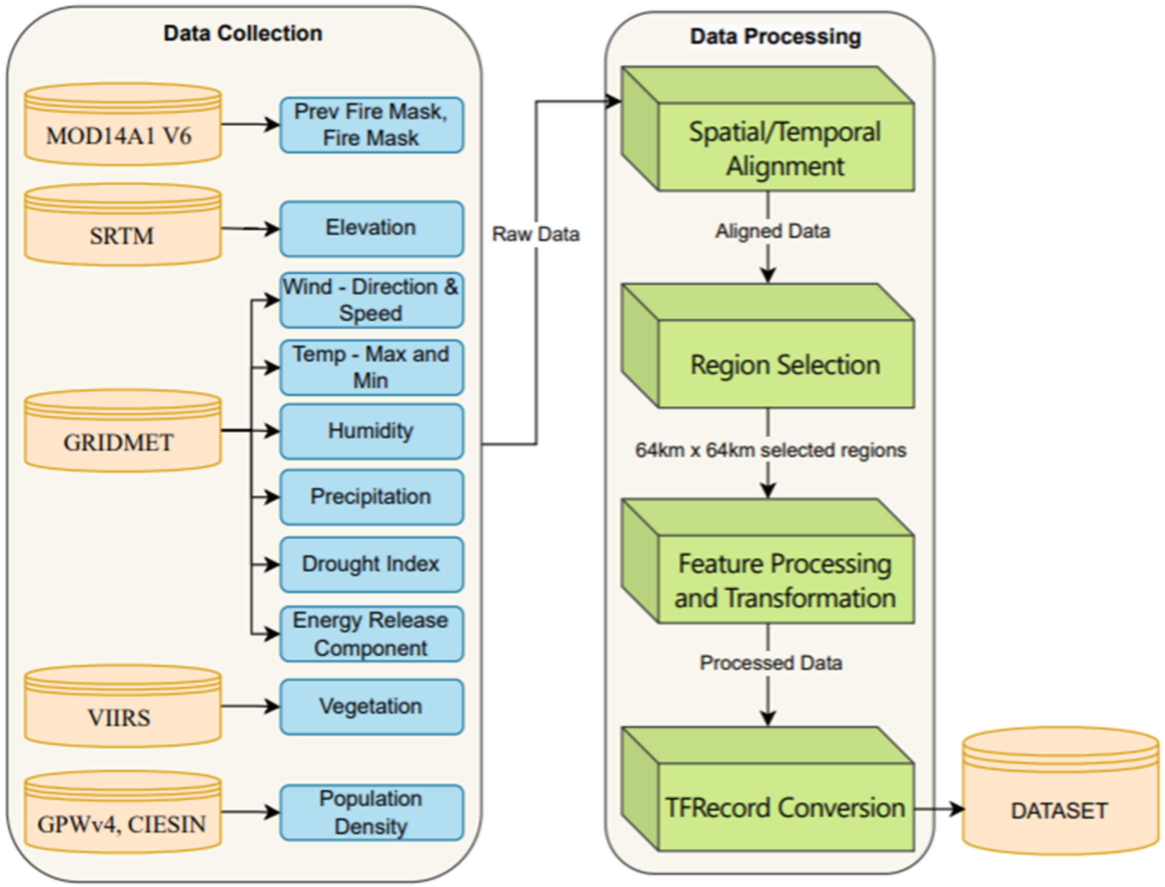

Spatial and Temporal Alignment

Resolution Adjustment

- Features resampled to a consistent 1 km resolution.
- Downsampling: High-resolution data (e.g., SRTM at 30 m).
- Upsampling: Low-resolution data (e.g., GRIDMET at 4 km, VIIRS at 0.5 km).

Temporal Aggregation

- Weather data recorded every 6 hours was averaged to daily metrics.

Dataset: Data Collection Framework

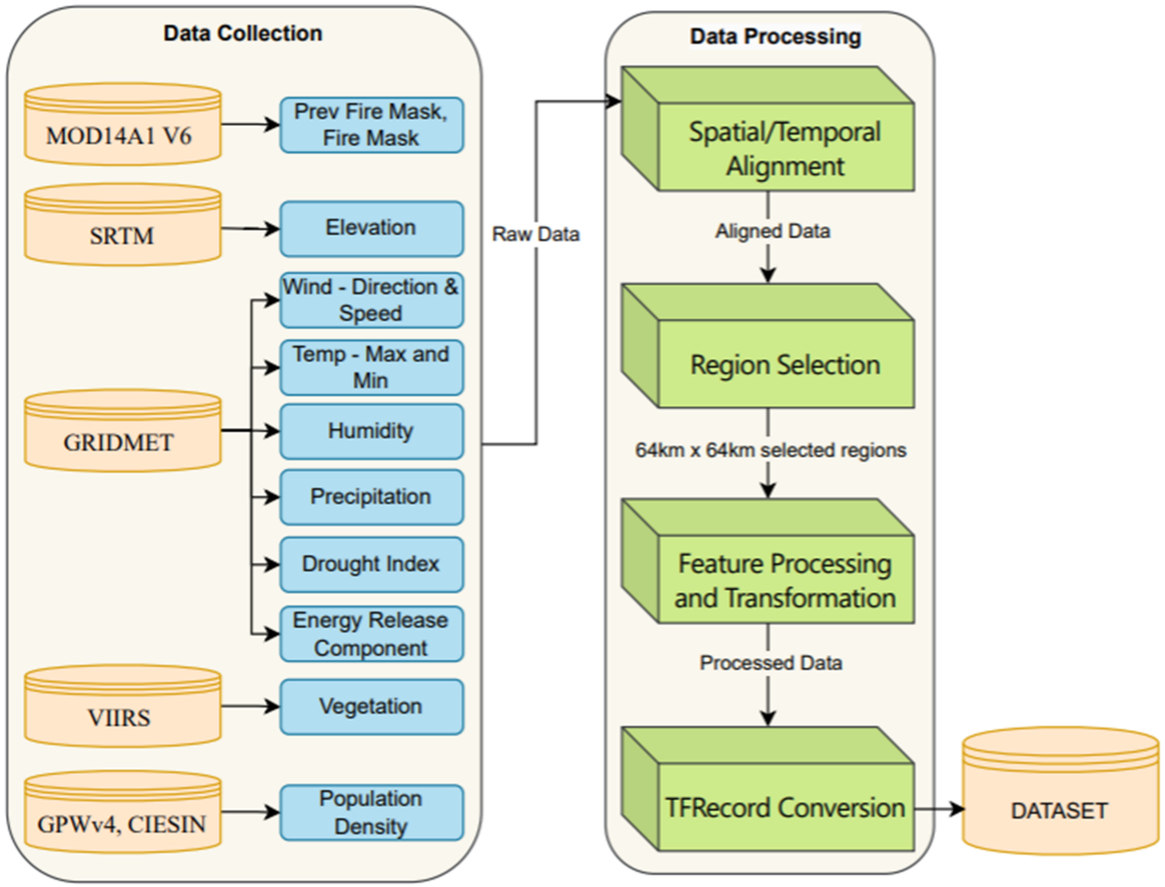

Region Selection

Active Fire Detection

- Focused on regions with detected fire activity on day t.
- Regions centered on 1 km × 1 km cells where fire was observed.

Region Size

- Each fire region expanded to a 64 km × 64 km area.
- Provides spatial context for fire spread while remaining computationally manageable.

Dataset: Data Collection Framework

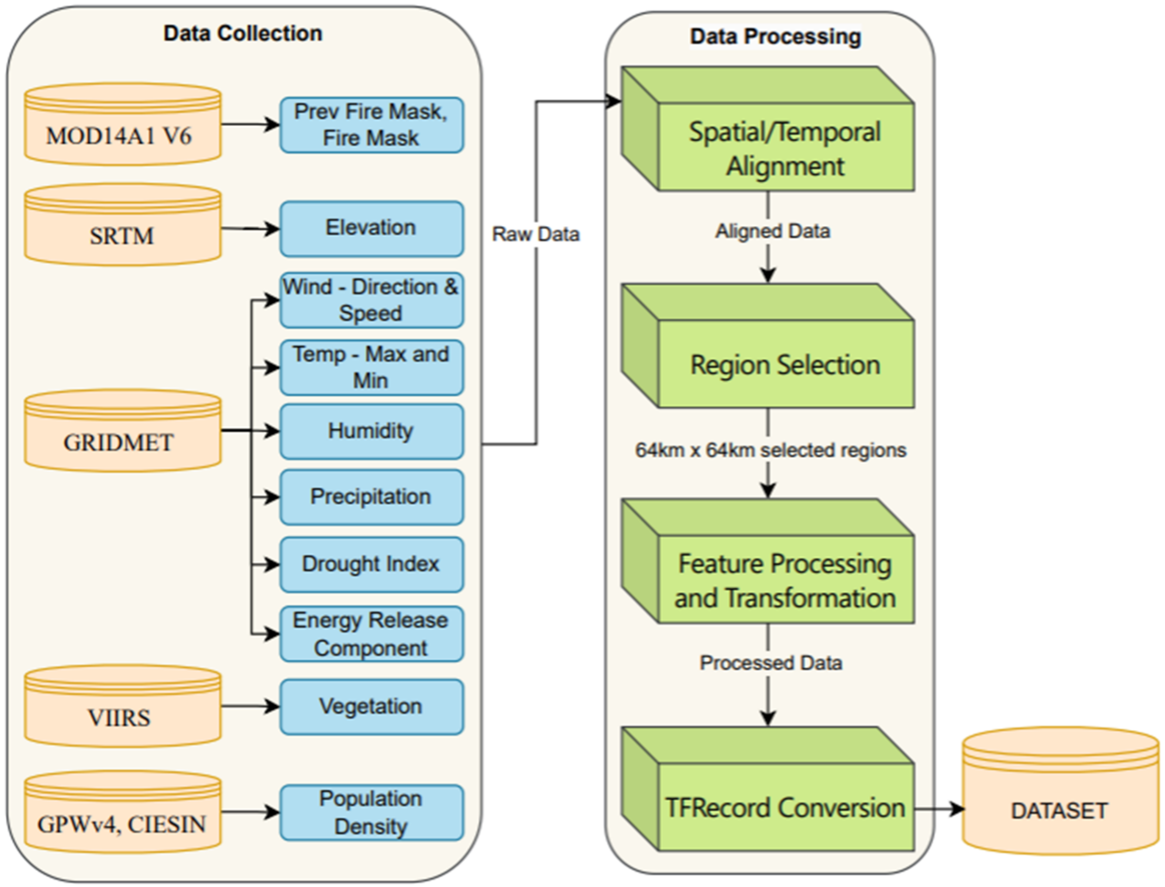

Feature Processing

Normalization and Clipping

- Features clipped to avoid extreme outliers:
  - Based on physically meaningful limits or data percentiles (0.1%–99.9%).
- Normalized by subtracting the mean and dividing by standard deviation.

Multi-Channel Image Format

- Features structured into a 64 × 64 grid, forming multi-channel images.
- Each channel represents a feature, with the next-day fire mask as the target label.

Dataset: Data Collection Framework

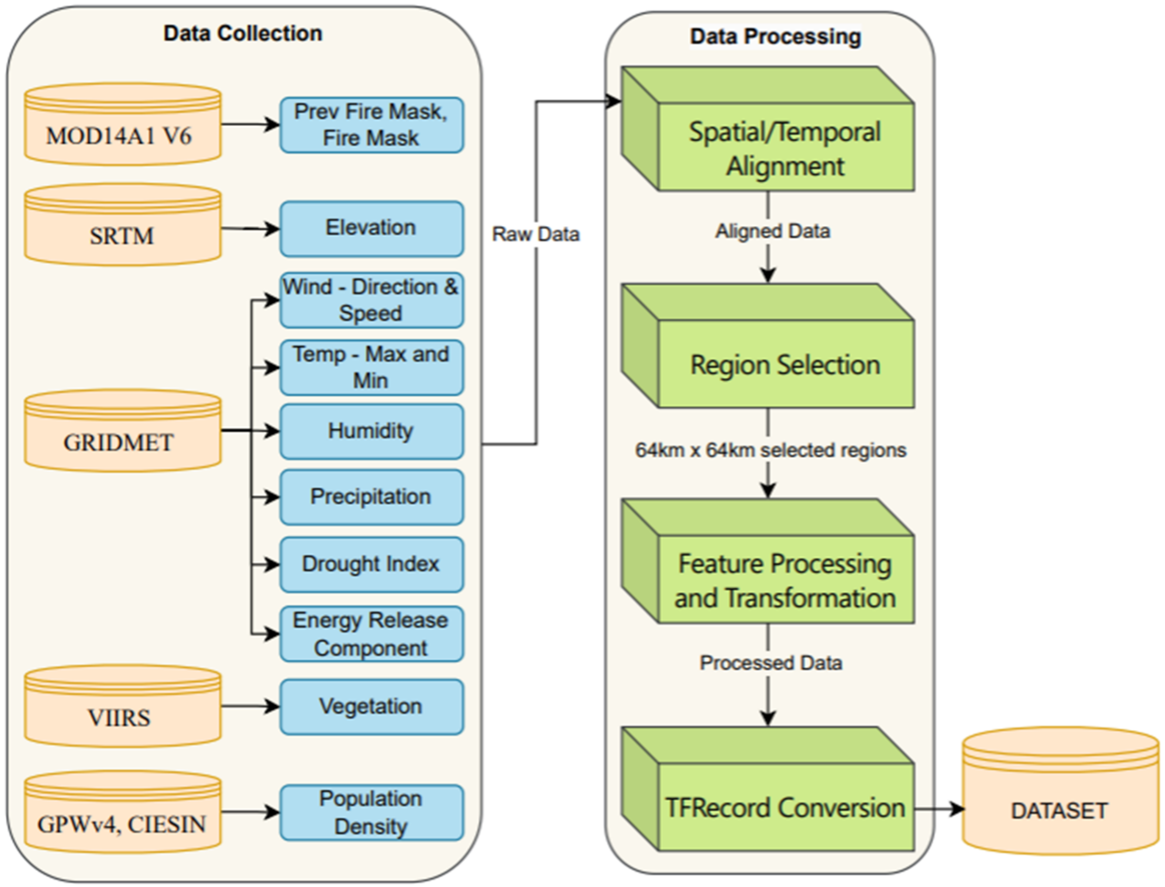

Data Storage and Format

TFRecord Conversion

- Dataset stored in TFRecord format, optimized for TensorFlow.
- Ensures efficient data access and loading for large-scale machine learning tasks.
- Each sample in the dataset represents a 64 km × 64 km grid with 1 km × 1 km cells, capturing spatial features around detected wildfires.

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Dataset: Features

Feature Category	Feature	Details
Topographical	Elevation	Measures the height of the land surface from sea level.
	Vegetation Cover	Quantifies the density and type of vegetation, which influences fuel for fires.
Anthropogenic	Population Density	Indicates the number of people living per unit area, crucial for evacuation and risk assessment.
Meteorological	Wind Direction	Shows the direction from which the wind is blowing, influencing fire spread direction.
	Wind Speed	Measures how fast the wind is blowing, a critical factor in fire spread rate.
	Minimum Temperature	Lowest daily temperature, affecting humidity and dryness of the area.
	Maximum Temperature	Highest daily temperature, affecting evaporation and dryness of the area.
	Humidity	Measures the amount of moisture in the air, impacting fire behavior.
	Precipitation	Amount of rainfall, which can reduce fire risk by moistening potential fuel.
Additional Variables	Drought Index	Indicates prolonged absence of precipitation, increasing the vulnerability of the area to fires.
	Energy Release Component (ERC)	Represents the potential heat energy release per unit area in the event of a fire.
	Fire Masks (time t and t+1)	Binary indicators showing the presence or absence of fire at two time points, essential for tracking fire progression.

Dataset: Preprocessing

- **Random cropping**: Avoid overfitting on centered fire regions.
- **Noise augmentation**: Improve robustness to data variability.
- **Filtering invalid data**: Remove uncertain labels (e.g., clouds).
- **Standardization**: Scale features to zero mean, unit variance.

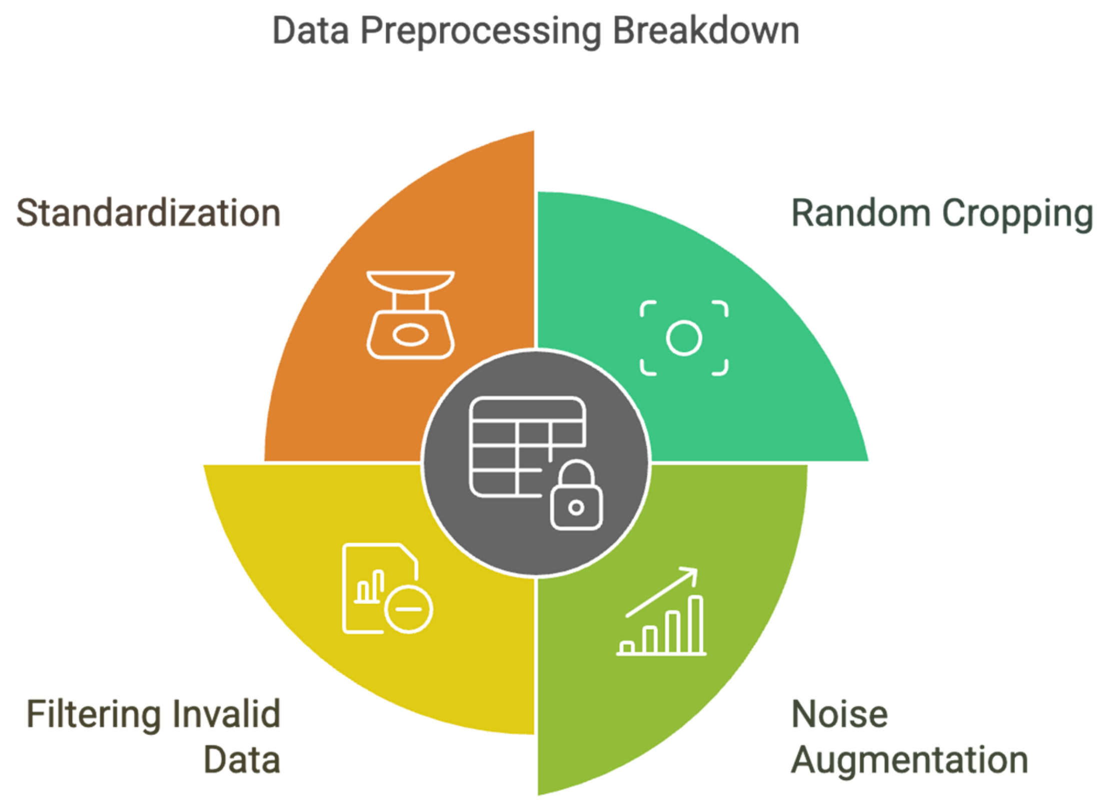

Dataset: Preprocessing

Random Cropping

- Applied to input grids to focus on different parts of fire regions.
- Reduces overfitting by preventing the model from memorizing centered fire areas.
- Ensures variability in training samples, improving generalization.

Filtering Invalid Data

- Removed data with missing or uncertain values (e.g., cloud interference).
- Ensures high-quality, reliable data for training.
- Prevents model degradation due to noisy or corrupted inputs.

Noise Augmentation

- Random noise added to features (e.g., wind speed, temperature).
- Mimics real-world variability in environmental data.
- Enhances robustness of the model to unpredictable inputs.

Standardization

- Features scaled to have **zero mean** and **unit variance**.
- Normalization aligns all data ranges, enabling efficient learning.
- Reduces the risk of one feature dominating due to larger magnitude.

Dataset: Example

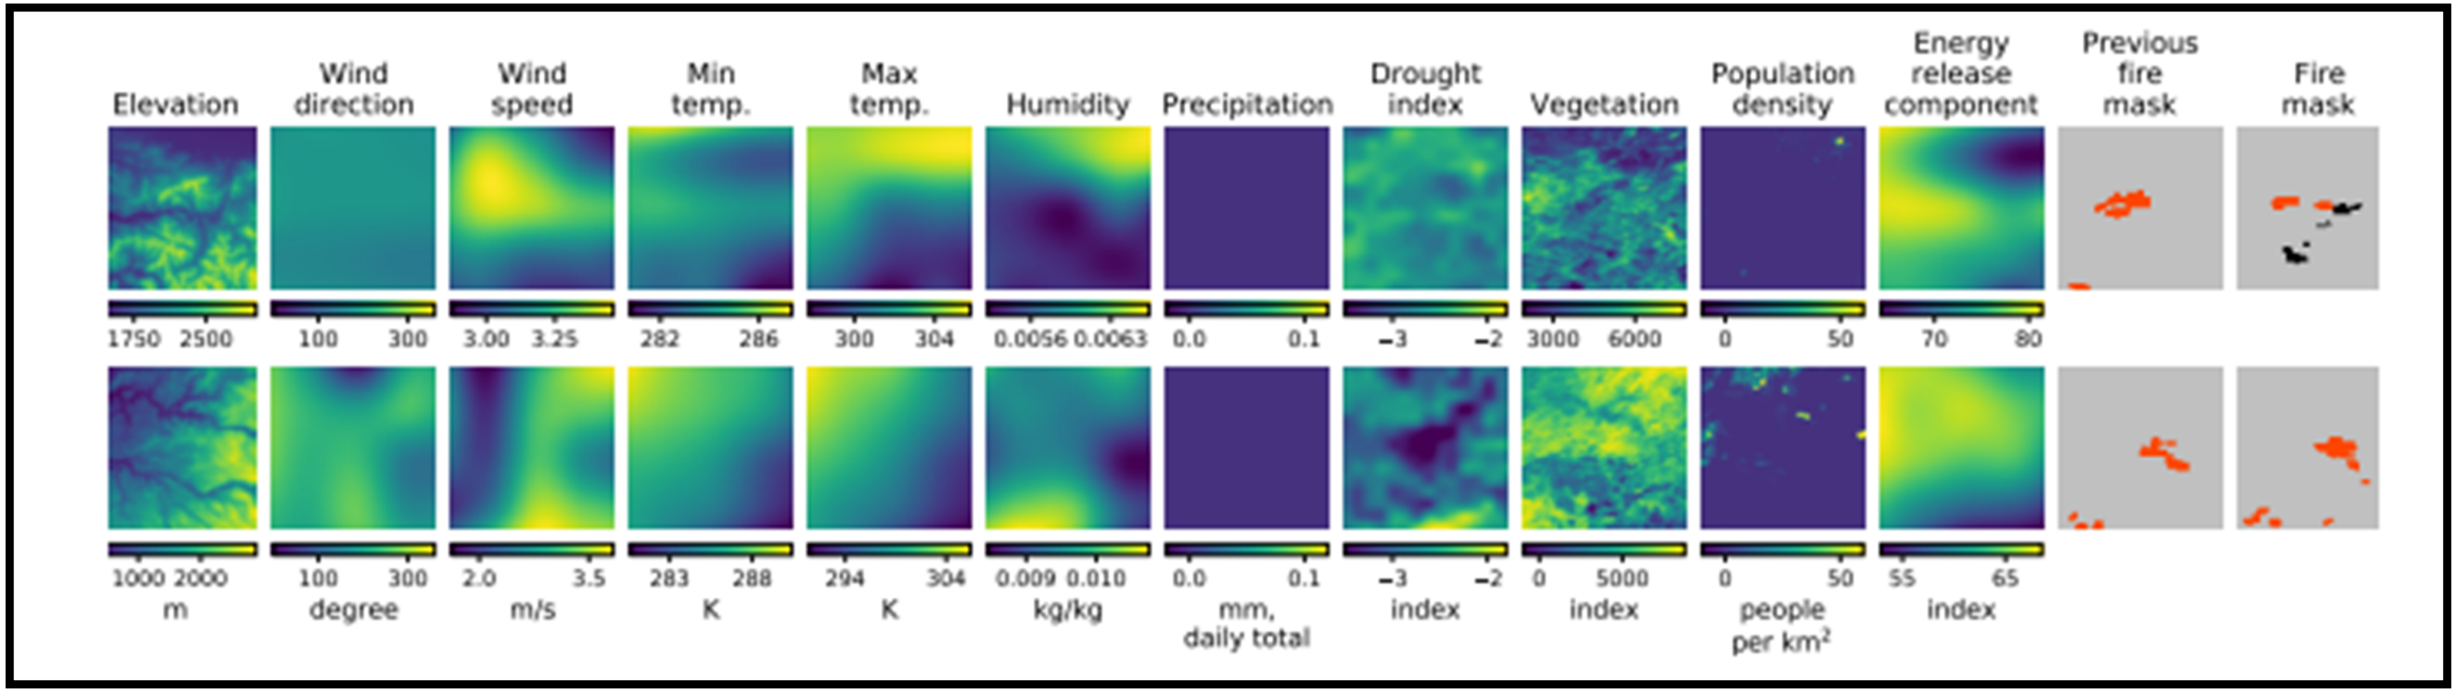

Dataset: Imbalance Ratio

- Fire regions <5% of total data, creating a major imbalance.
- Risk of bias toward "no fire" predictions.
- Advanced loss functions (Focal Tversky Loss) are employed to mitigate this issue.

Datasets	Fire (%)	No Fire (%)
Training Set	3.60	96.40
Validation Set	4.17	95.83
Test Set	3.74	96.26

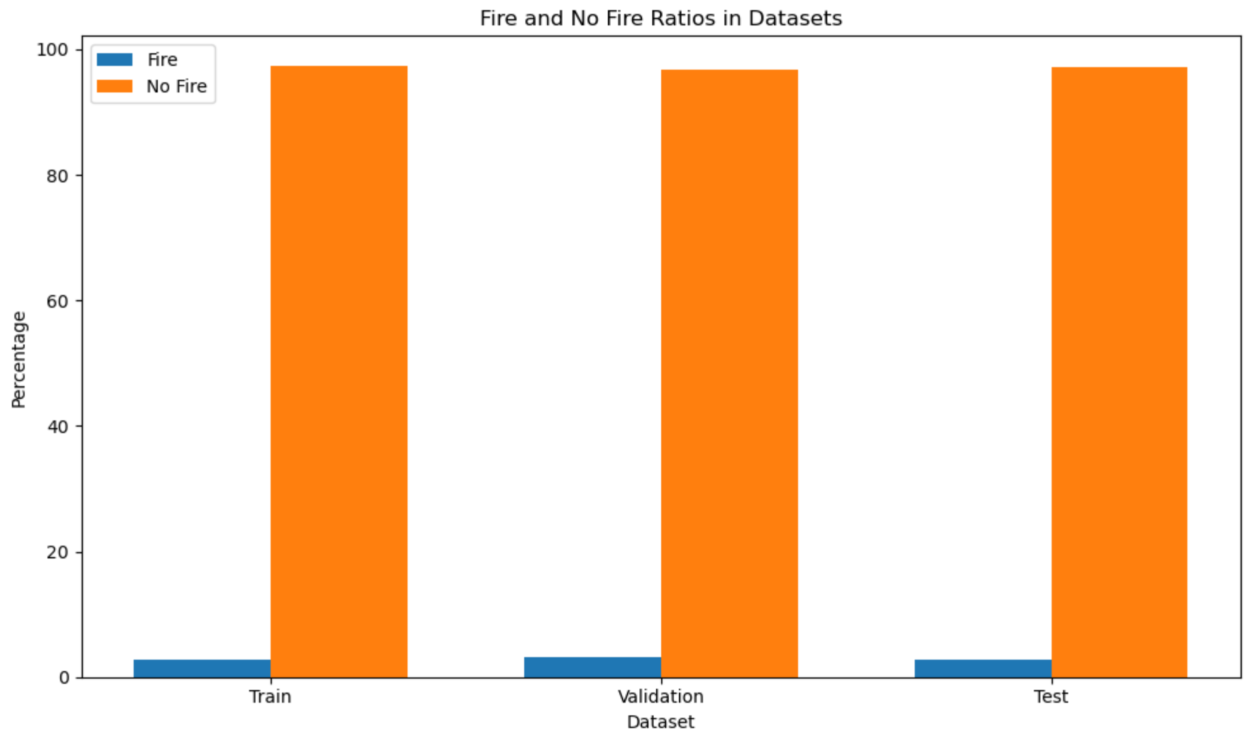

Dataset: Addressing Data Imbalance

Why Data Imbalance is Critical

- Wildfire datasets are **highly imbalanced**, with fire pixels often making up less than **5%** of the data.
- Models trained on such data are prone to favor the majority class ("no fire"), leading to poor sensitivity for detecting fire regions.
- A **high false-negative rate** in fire prediction could have severe real-world consequences, such as missed fire zones during emergencies.

Methodology

Methodology: Wildfire Prediction Workflow

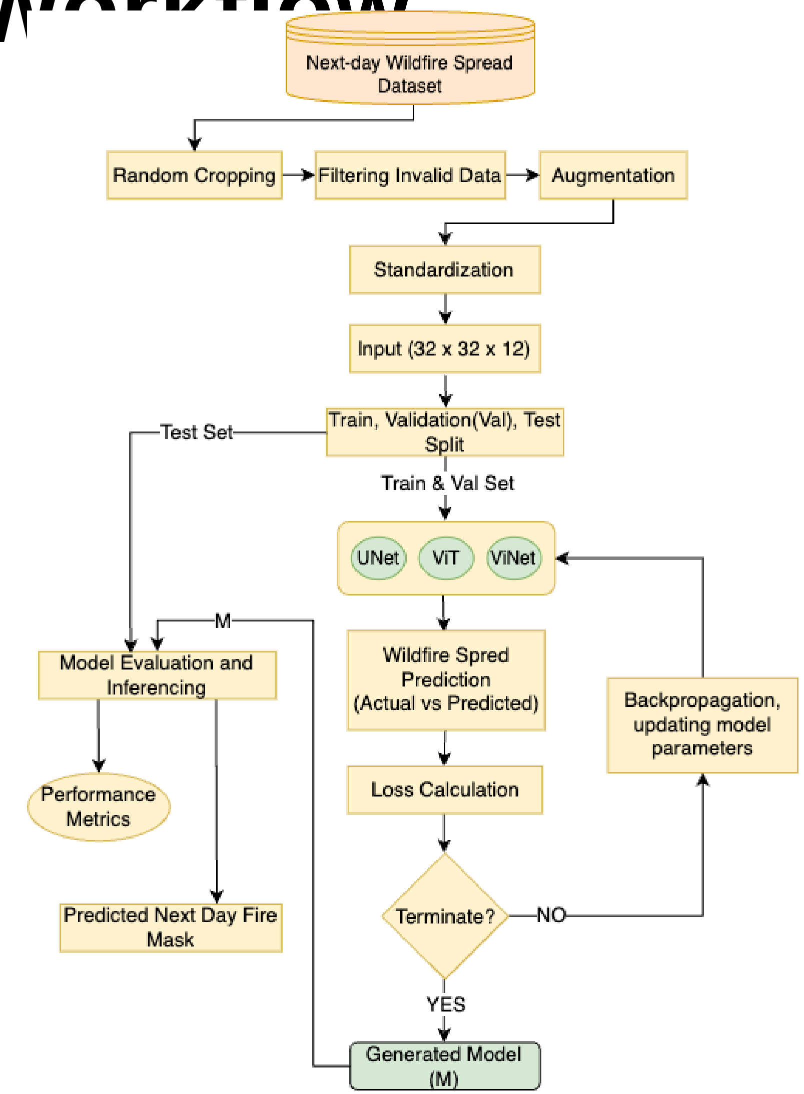

Key Workflow Components:

1. Data Processing
   - Input: 32x32x12 format
   - Steps: Cropping → Filtering → Augmentation → Standardization
2. Model Pipeline
   - Three parallel models: UNet, ViT, ViNet
   - Train/Val/Test split
   - Loss calculation and backpropagation loop
3. Output Generation
   - Performance metrics are generated to assess accuracy
   - Final output is a predicted next-day fire mask
   - Generated model (M) is produced after successful training

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: Input & Output Structure

- **HWC Format:** In the context of our image data, HWC stands for Height, Width, and Channels. This format is used to represent images where:
  - H: Height of the image (number of rows).
  - W: Width of the image (number of columns).
  - C: Number of channels (each channel represents a different feature).

- Input Tensor:
  - Shape: [H, W, C]
  - Example Shape: [64, 64, 12]

  (for a 64x64 image with 12 input features)

- Output Tensor:
  - Shape: [H, W, C]
  - Example Shape: [32, 32, 1]

  (for a 32x32 output image with 1 output feature, the predicted fire mask)

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: U-Net Architecture

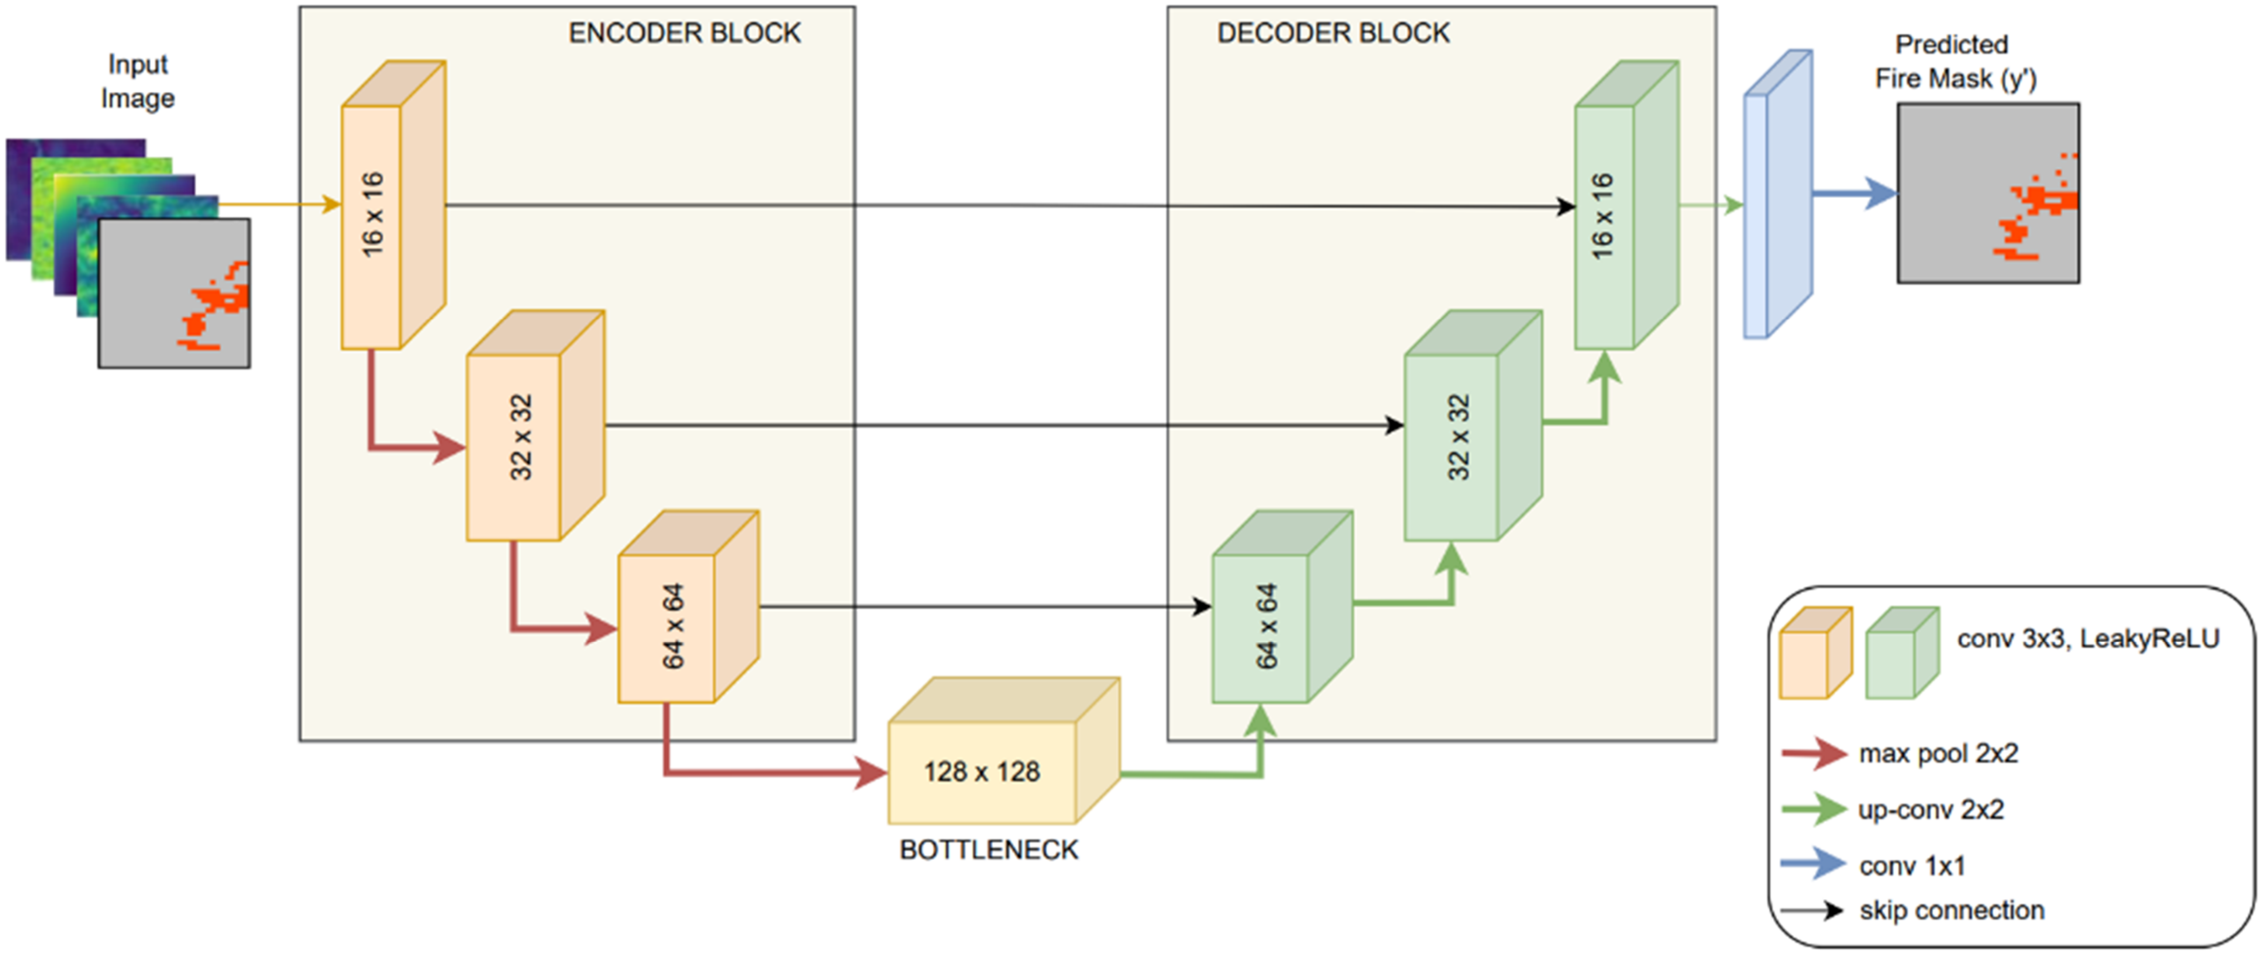

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: U-Net Model Architecture

Encoder-Decoder Structure:

- Encoder: Extracts hierarchical spatial features.
- Decoder: Reconstructs high-resolution fire masks from encoded features.
- Skip connections between encoder and decoder preserve fine-grained details.

Relevance to Wildfire Prediction:

- Excels in segmenting fine-scale fire regions essential for mitigation planning.

Specialization:

- Tailored for pixel-level segmentation tasks like wildfire boundary detection.
- Localized feature extraction ideal for identifying small fire regions.

Strengths:

- High precision in capturing detailed fire boundaries.
- Efficient processing of high-resolution spatial grids.

Methodology: U-Net Limitations and Challenges

Focus on Local Features:

- Struggles to capture long-range spatial dependencies.
- Limited contextual understanding for large-scale fire propagation.

Complex Interactions:

- Unable to model interactions between distant fire-prone areas.
- Loses global context critical for wildfire spread prediction.

Scalability Issues:

- Performance degrades with increasing region sizes.
- Insufficient when processing multimodal spatial-temporal data.

Summary:

- U-Net excels in localized segmentation but lacks the broader spatial context needed for wildfire prediction.

Methodology: Vision Transformers(ViT) Architecture

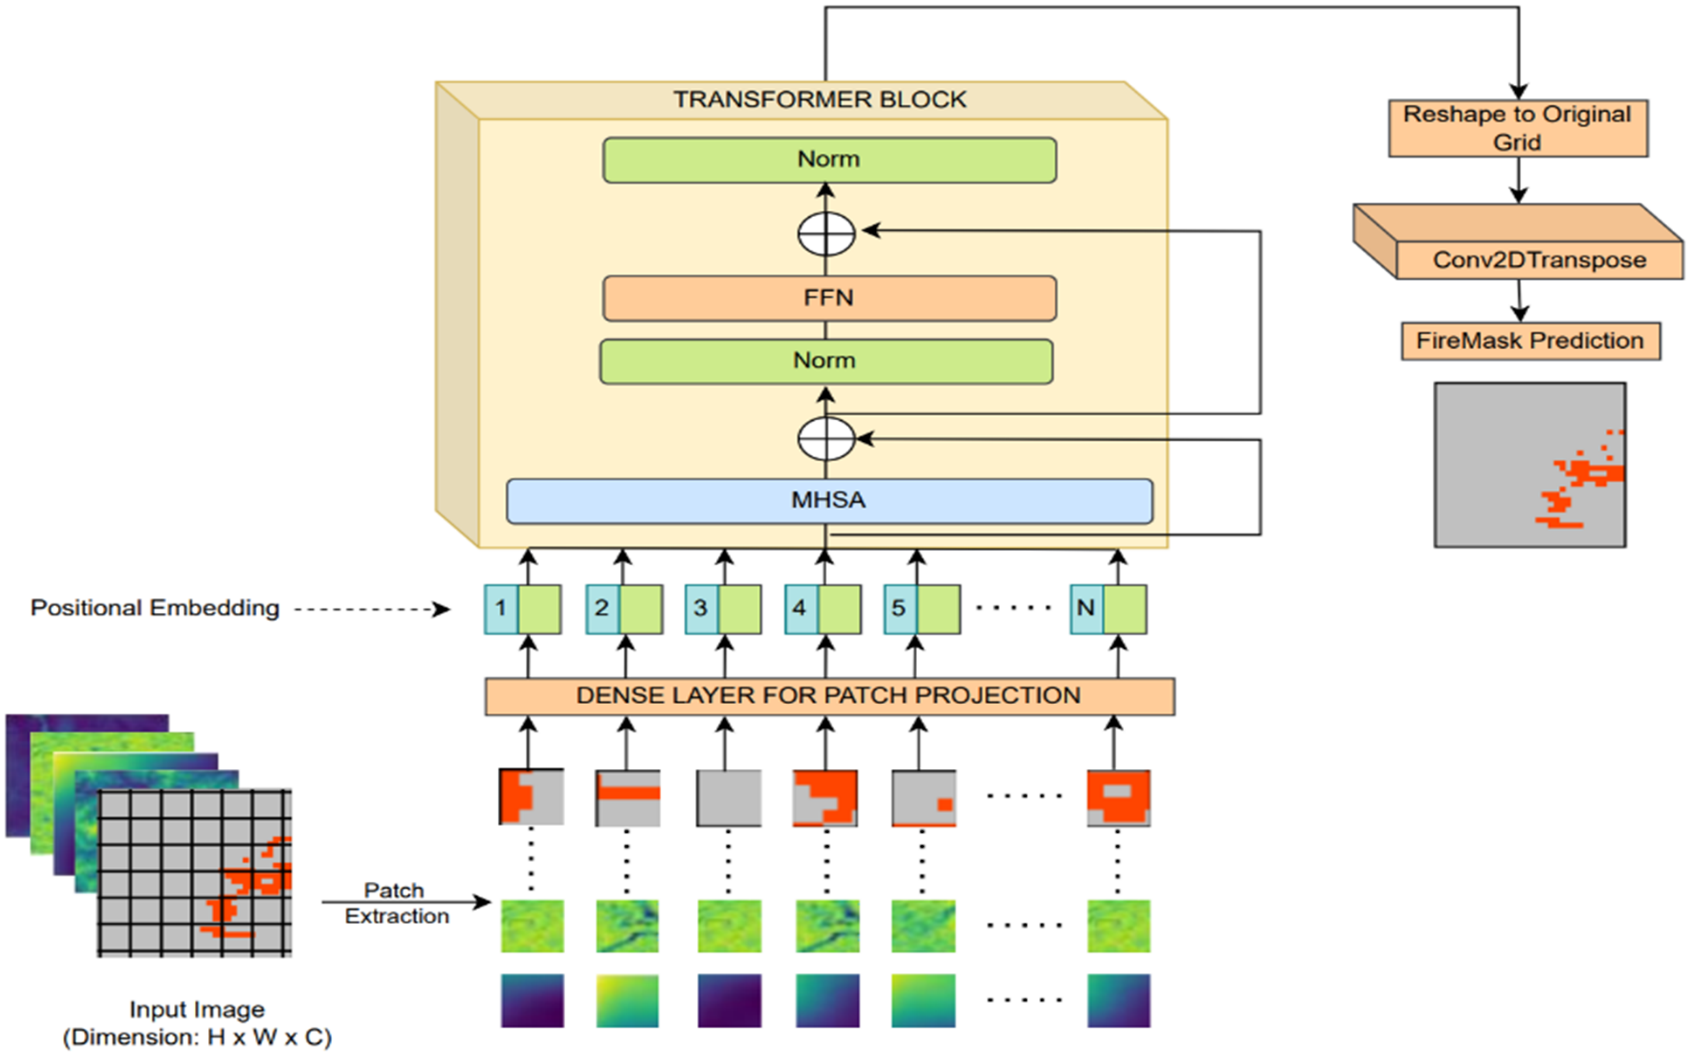

Methodology: Vision Transformers(ViT) Architecture

Self-Attention Mechanism:

- Captures long-range dependencies across spatial regions.
- Analyzes relationships between all pixels simultaneously.

Patch-Based Processing:

- Divides input grid into smaller patches.
- Embeds patches as feature vectors for global analysis.

Global Context:

- Excels in understanding large-scale patterns in fire dynamics.
- Ideal for modeling wildfire propagation across vast areas.

Flexibility:

- Adapts easily to diverse spatial resolutions and multimodal inputs.

Methodology: ViT Limitations and Challenges

Lack of Fine-Grained Precision:

- Struggles to capture small-scale, localized details in fire regions.
- Overshadowed by U-Net in precise boundary detection.

Computational Demands:

- High memory and processing requirements for large input grids.
- Training time increases with larger datasets and grid resolutions.

Overfitting Risks:

- Requires extensive data for effective training.
- Sensitive to noise in input features, reducing robustness.

Summary:

- ViT provides global insights but lacks fine-grained segmentation capabilities, making it insufficient for precise wildfire boundary prediction.

Methodology: Vi-Net Architecture

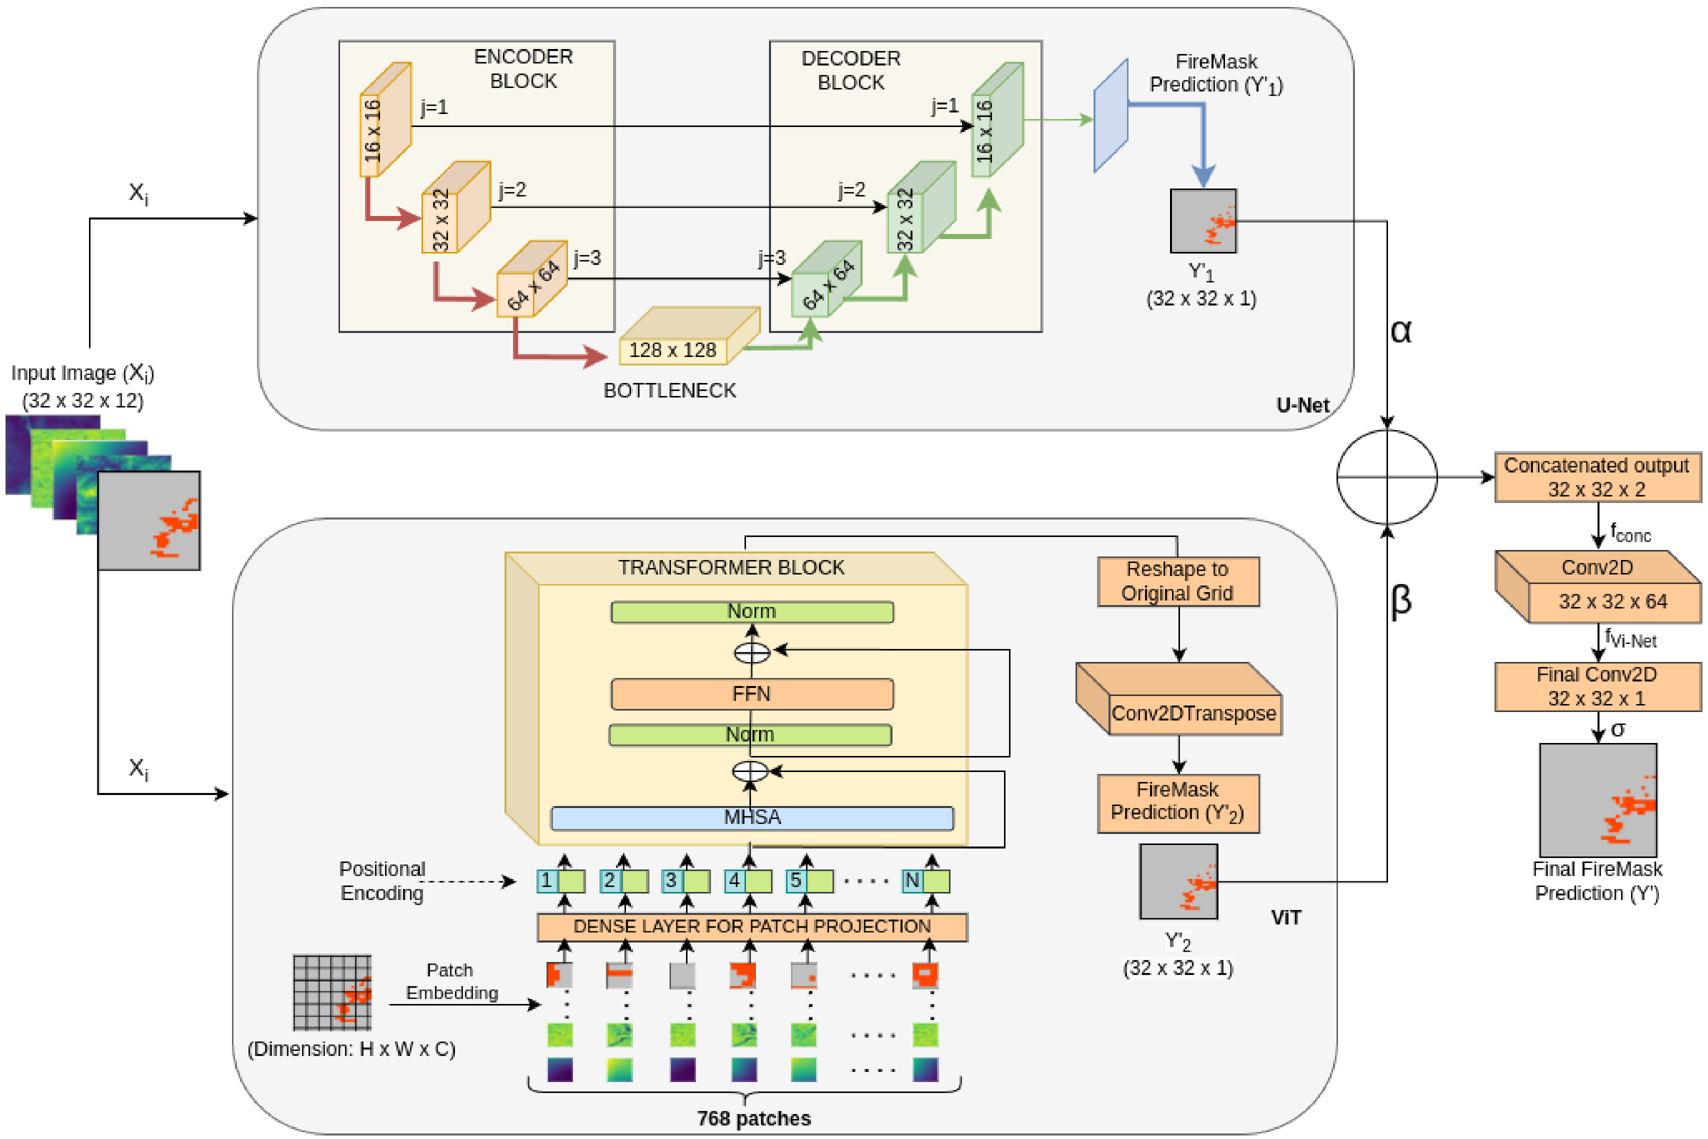

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: Vi-Net Architecture

Hybrid Approach:

- Combines outputs of U-Net and ViT architectures.
- Balances fine-grained segmentation with global context understanding.

Workflow:

- U-Net extracts localized spatial details.
- ViT models long-range dependencies across the input grid.
- Outputs from both models are merged to generate final predictions.

Core Components:

- Encoder (U-Net): Captures precise fire region boundaries.
- Global Context Module (ViT): Integrates large-scale fire spread patterns.

Methodology: Vi-Net Architecture

Fusion Process:

- Outputs of U-Net and ViT combined through weighted summation: $F_{\text{Vi-Net}} = \alpha \times f_{\text{U-Net}} + \beta \times f_{\text{ViT}}$
- Parameters α and β adjust the contributions of each model.

Integration Logic:

- U-Net provides high-resolution fire masks.
- ViT contextualizes these masks with global insights.

Final Output:

- Enhanced binary segmentation mask representing next-day fire spread.

Advantages of Fusion:

- Retains U-Net's precision and ViT's contextual awareness.
- Produces robust, balanced predictions for wildfire management.

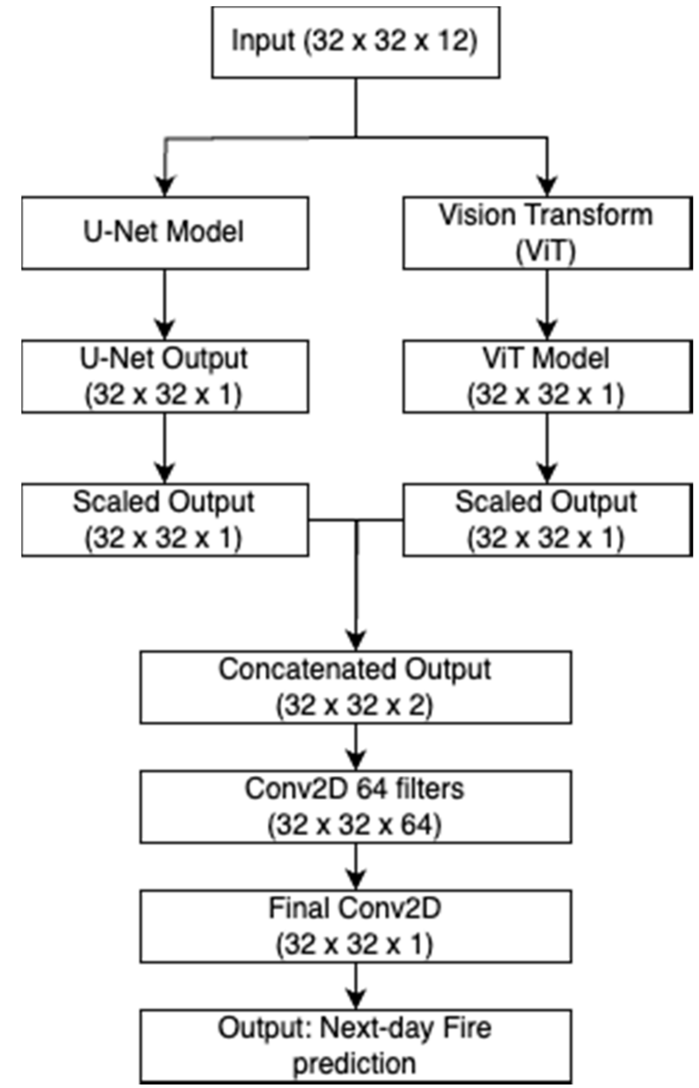

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: Vi-Net Strengths

Overcoming U-Net Limitations:

- Adds global spatial awareness to U-Net's localized predictions.
- Reduces the risk of underpredicting large-scale fire spread.

Robust Hybrid Architecture:

- Provides the best of both worlds: fine-grained segmentation and global insights.
- Balances sensitivity (recall) and specificity (precision) for accurate predictions.

Addressing ViT Shortcomings:

- Incorporates ViT's global context into precise fire region boundaries.
- Improves performance on smaller, localized fire zones.

Scalability and Adaptability:

- Adapts to diverse fire scenarios and multimodal input data.
- Effective for real-world wildfire prediction and mitigation planning.

Methodology: Path Planning Module

- **A* Search** algorithm, unlike other traversal techniques, has “intuition”.
- Like Dijkstra, A* works by making a lowest-cost path tree from the start node to the target node. However, the A* algorithm introduces a heuristic into a regular graph-searching algorithm, essentially planning ahead at each step so a more optimal decision is

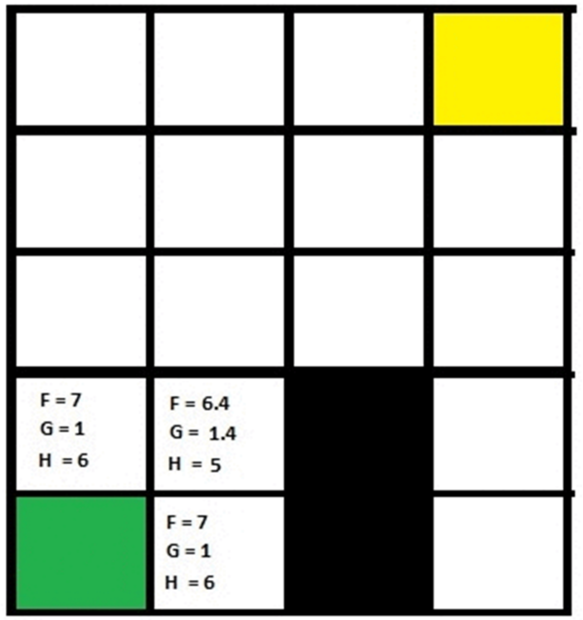

- A* expands paths that are already less expensive by using this function: $f(n)=g(n)+h(n)$,

Where:

- $f(n)$ = total estimated cost of path through node n
- $g(n)$ = cost so far to reach node n
- $h(n)$ = estimated cost from n to goal. This is the heuristic part of the cost function.

Methodology: Path Planning Module

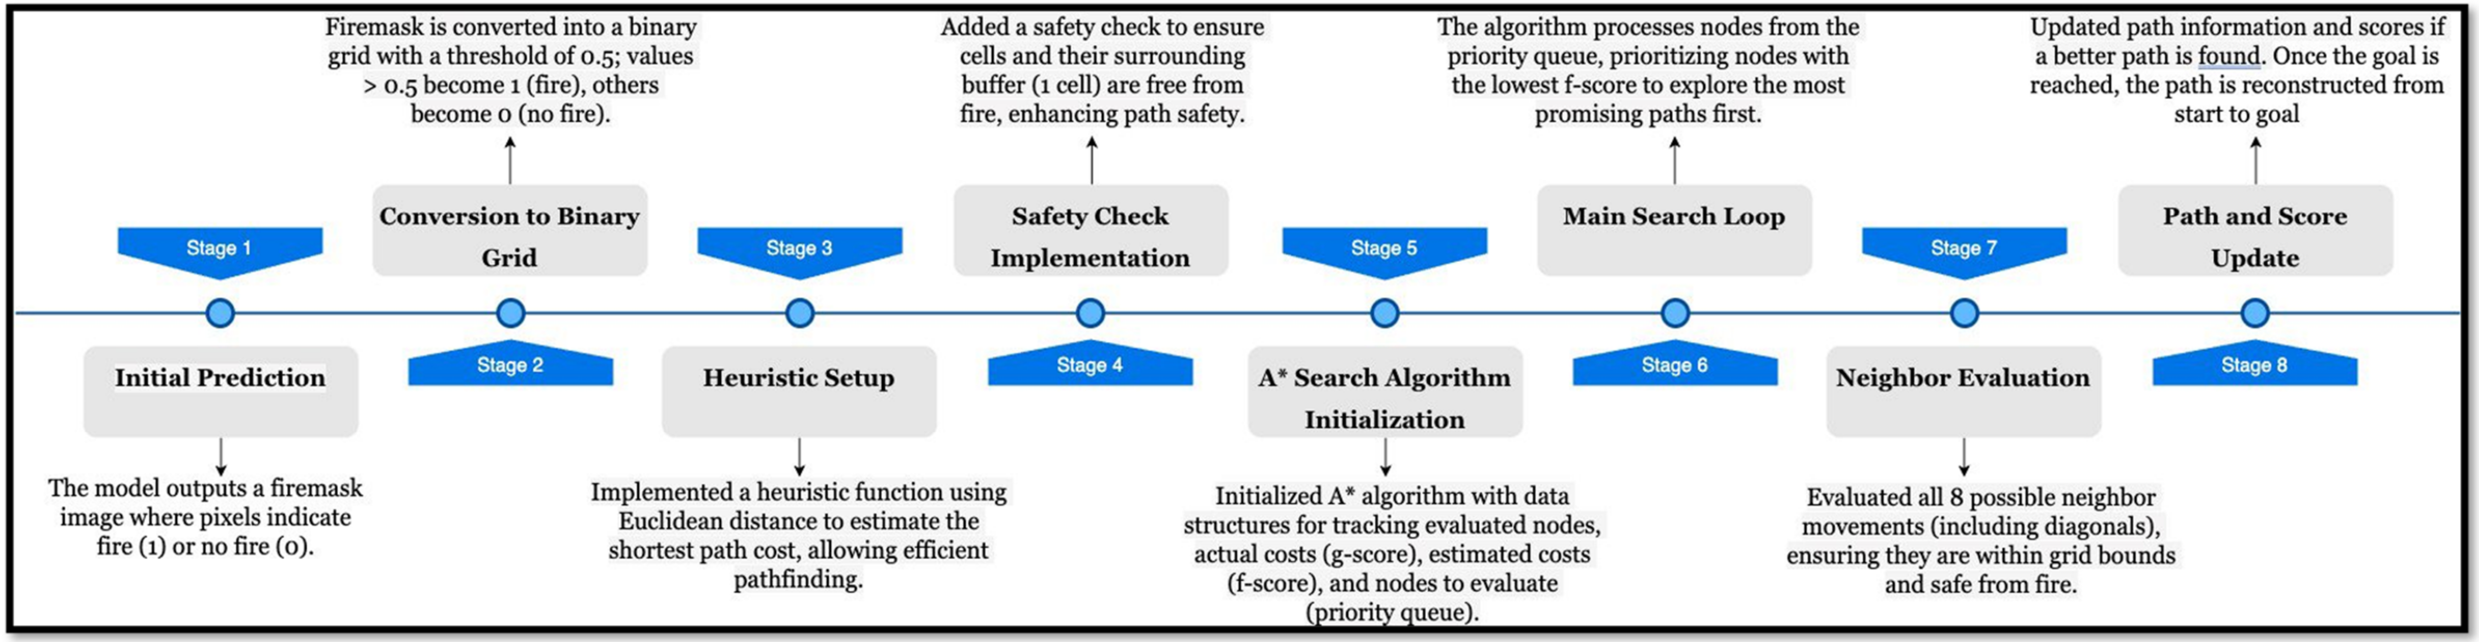

Methodology: Safe Path Planning with A*

Dynamic Fire Spread Prediction

- The algorithm integrates real-time fire spread predictions from Vi-Net to navigate changing conditions effectively.
- Adapts to fire progression, ensuring safer routing even in dynamic environments.

Modified A* Algorithm

- The A* algorithm calculates the shortest possible path from the start node to the goal node while avoiding hazardous areas.
- Incorporates fire avoidance heuristics to prioritize safety over speed.

Methodology: OA* Algorithm

Fire-Prone Regions as Obstacles

- Nodes marked as 1 (fire-affected) or within a defined buffer zone are considered impassable.
- Nodes in the grid labeled as 0 are for safe navigable regions.

Buffer Zone Implementation

- Adds an adjustable safety margin around fire-affected cells.
- This ensures no path is planned too close to fire zones, accounting for risks like heat or smoke.

Efficiency Through Heuristics

- Utilizes the Euclidean distance heuristic to calculate costs, ensuring computational efficiency without compromising accuracy.
- Evaluates eight possible movement directions (including diagonals) to allow flexible and optimal pathfinding.

Shortest and Safest Route Generation

- Balances the dual objectives of minimizing distance and maximizing safety.
- Ensures paths are computationally feasible in real time, critical for emergency scenarios.

Methodology: Vi-Net and Path Planning Integration

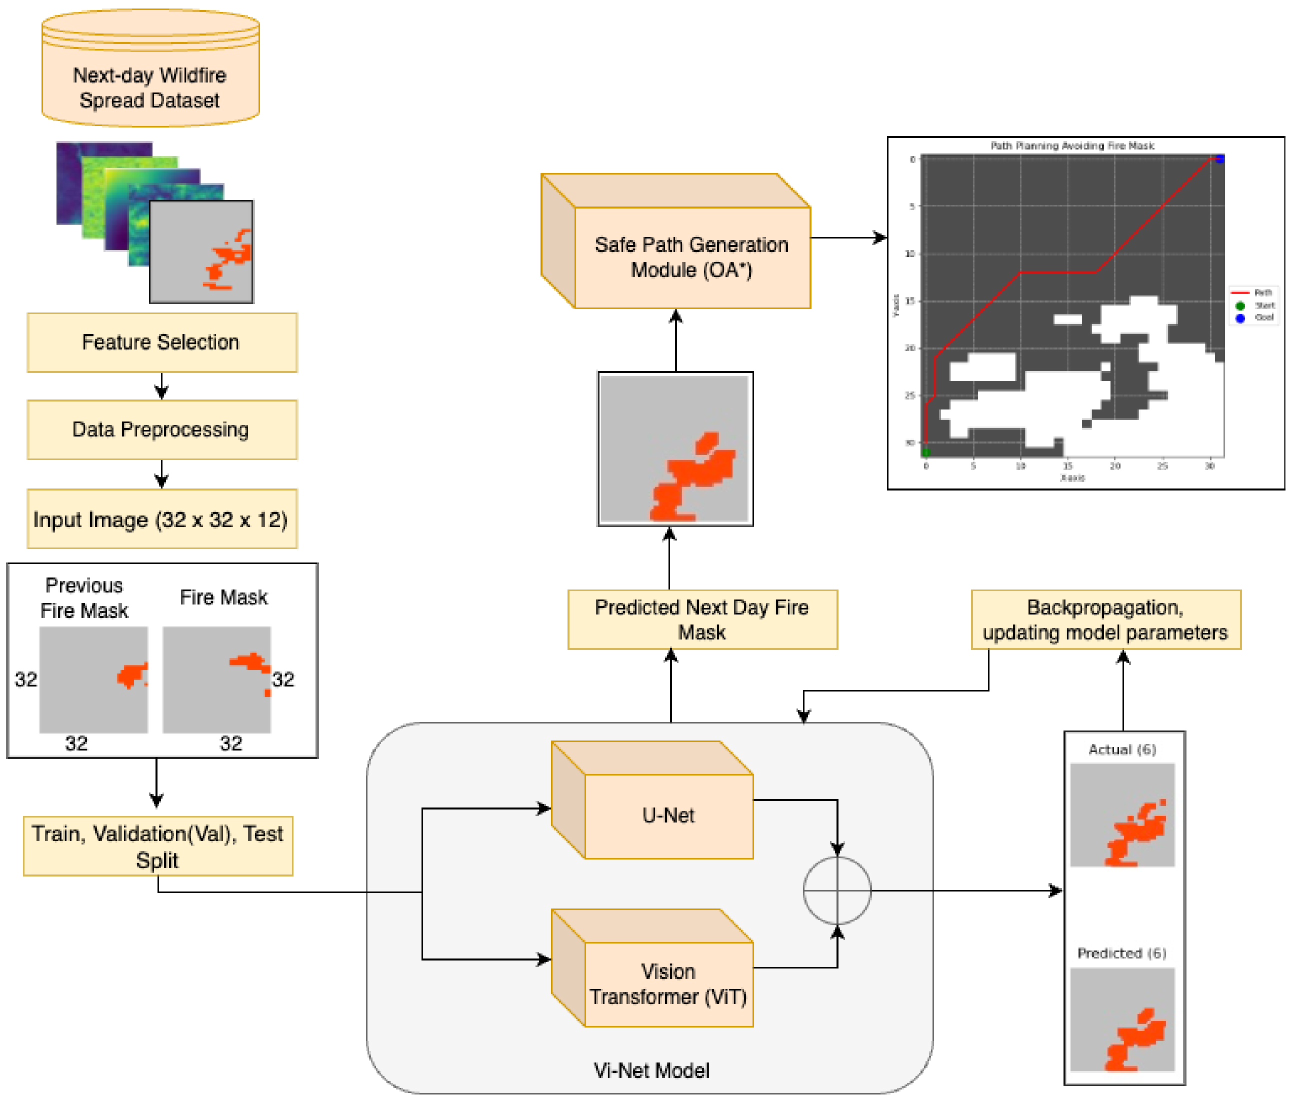

- Vi-Net outputs used directly in A* algorithm.
- Enables actionable insights for emergency responders.
- Seamless workflow for prediction and mitigation.

<u>End-to-End System:</u>
Predict → Avoid → Navigate

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: Vi-Net and Path Planning Integration

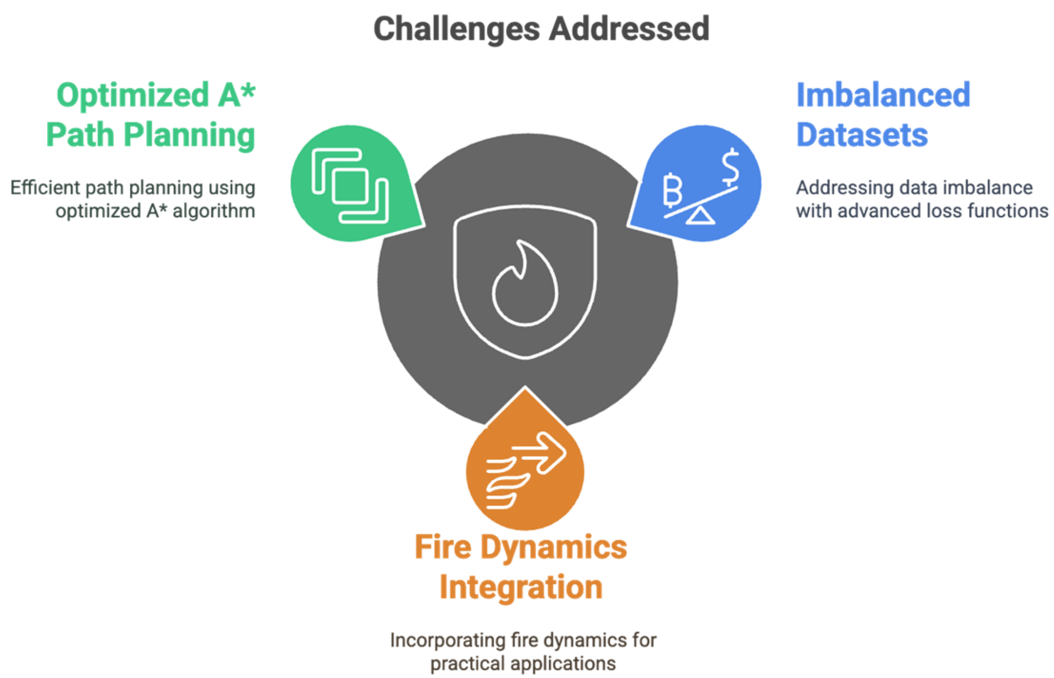

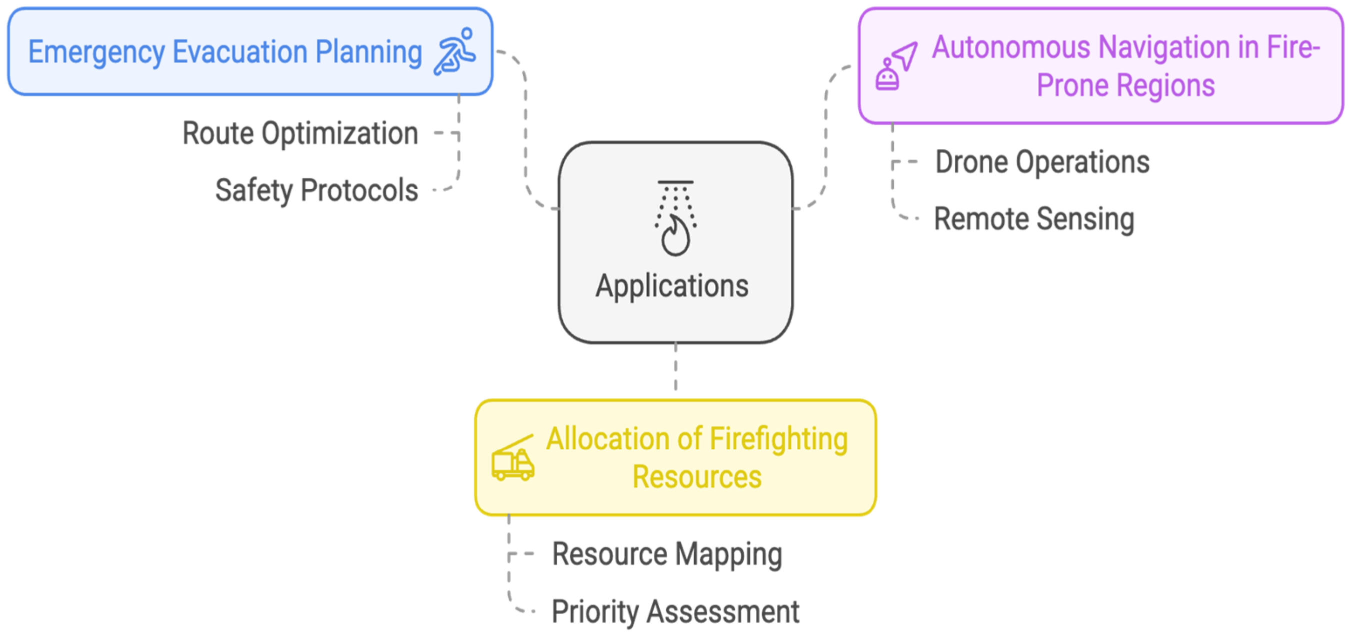

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Methodology: Performance Metrics

Term	**Definition**
True Positives (TP)	Regions correctly predicted as fire by the model that are actually on fire in the ground truth.
True Negatives (TN)	Regions correctly predicted as non-fire when they are indeed not on fire in the ground truth.
False Positives (FP)	Regions incorrectly predicted as fire by the model that are not on fire in the ground truth, leading to overprediction.
False Negatives (FN)	Regions incorrectly predicted as non-fire by the model that are actually on fire in the ground truth, leading to underprediction.

Methodology: Performance Metrics

Term	Definition
Precision	Precision measures the percentage of areas the model marked as fire that are actually on fire in reality. Particularly valuable in scenarios where the cost of a false positive is high.
Recall	Recall measures the percentage of real fire regions that the model correctly identified as fire. Critical in situations where missing a positive instance is costly.
F1 Score	F1 Score is the harmonic mean of precision and recall. It is used to balance the trade-offs between precision and recall in a single metric, which is especially useful when dealing with imbalanced datasets.
Jaccard Index(IoU)	The Jaccard index, also called Intersection over Union(IoU), is a measure of the similarity between two sets. Evaluates overlap between regions of predicted and ground truth sets.

Results and Experiments

Results: Experimental Setup

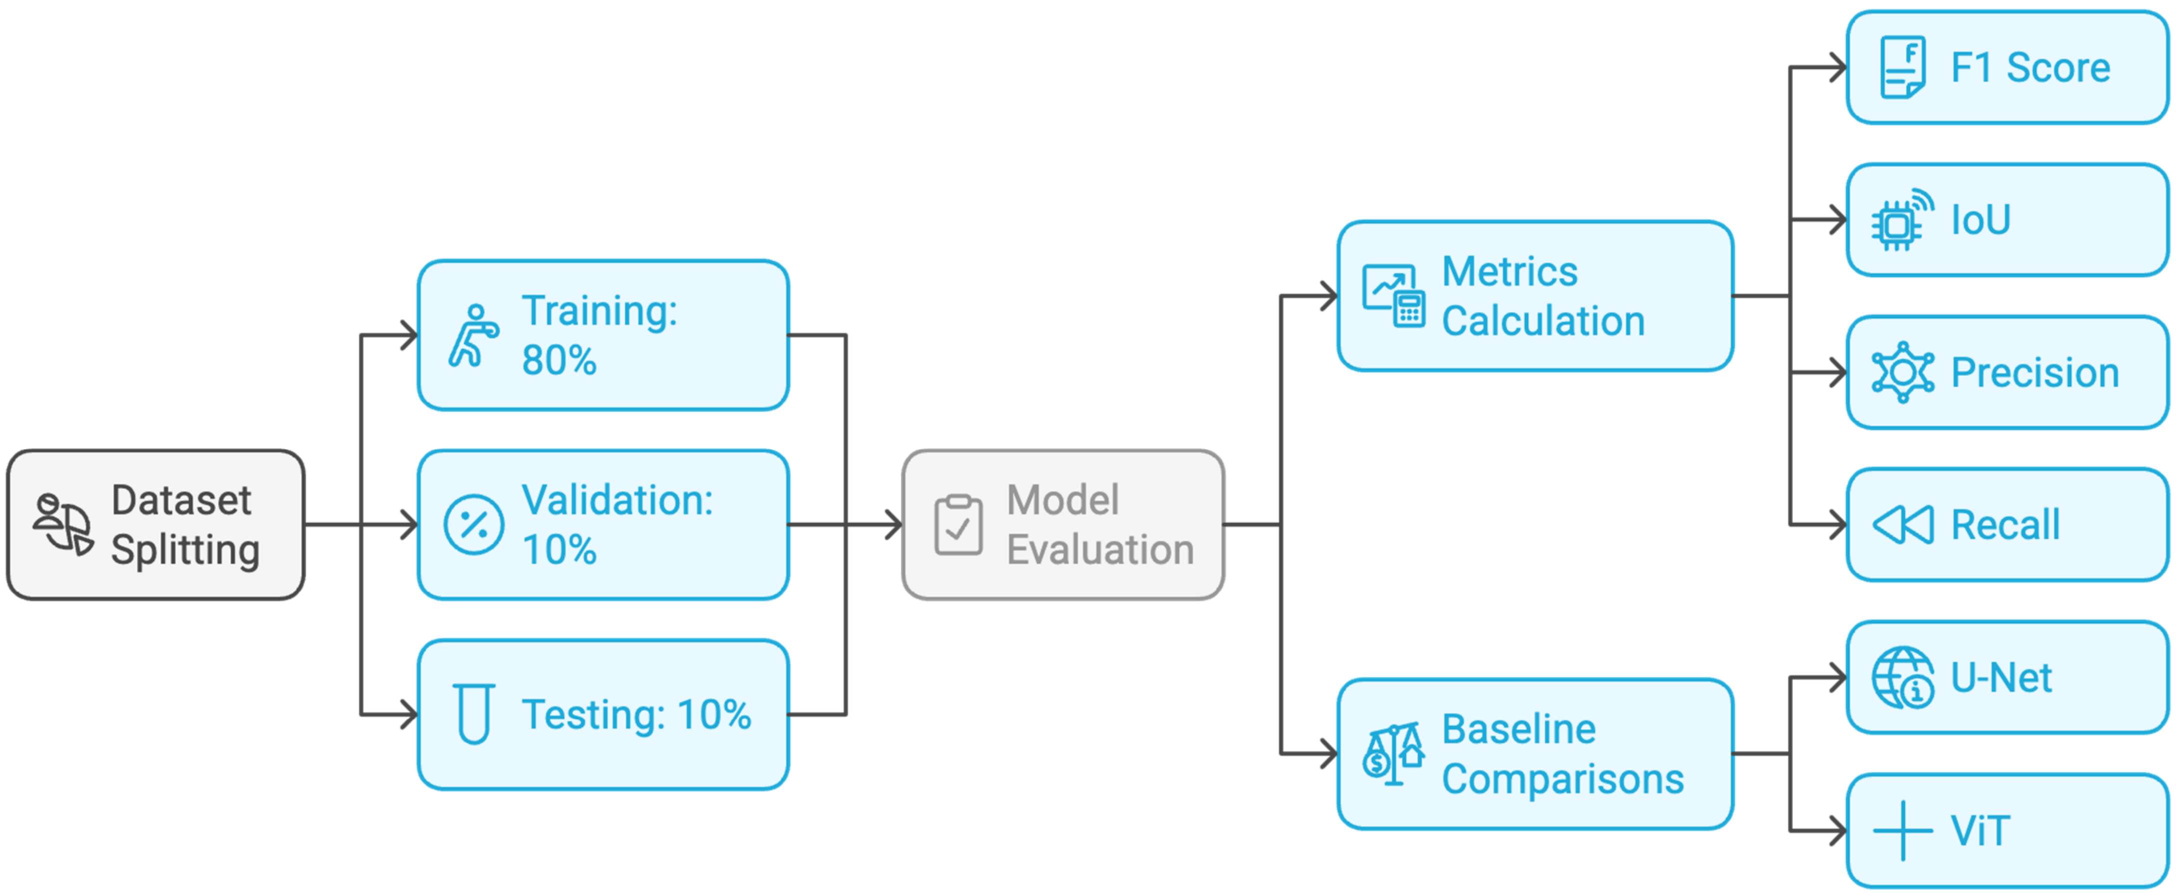

Results: Loss Functions

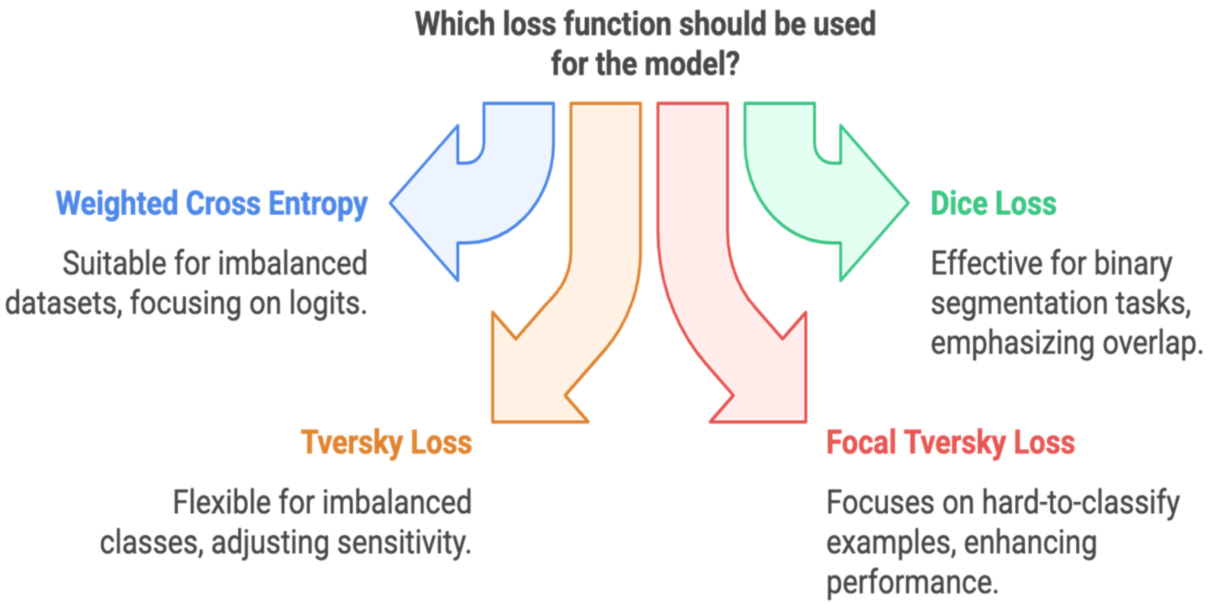

Models	Loss Functions	Precision	Recall	F1 Score
U-Net	WCE	0.9433	0.7018	0.8048
	Dice Loss	0.7491	0.8712	0.8056
	Tversky Loss	0.9957	0.7050	0.8255
	Focal Tversky Loss	0.9542	0.7506	0.8403
ViT	WCE	0.8492	0.5671	0.6801
	Dice Loss	0.9815	0.5635	0.7159
	Tversky Loss	0.9939	0.5559	0.7130
	Focal Tversky Loss	0.9873	0.5941	0.7418
Vi-Net	WCE	0.9250	0.9322	0.9286
	Dice Loss	0.9656	0.8999	0.9316
	Tversky Loss	0.9626	0.9418	0.9521
	Focal Tversky Loss	0.9834	0.9619	0.9725

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Results: Challenges of Loss Functions

Binary Cross Entropy (BCE):

- Treats all classes equally, failing to address imbalanced datasets.
- Results in models over-predicting "no fire" regions.

Dice Loss:

- Improves on BCE by focusing on overlaps between predicted and actual fire regions.
- However, it still struggles to prioritize minority fire pixels sufficiently.

Tversky Loss:

- Tversky Loss is a generalized version of Dice loss, incorporating a weighting mechanism to balance the false positives (FP) and false negatives (FN) during training.

Results: Loss Functions

Focal Tversky Loss

- A specialized loss function designed to prioritize fire regions.
- Combines concepts from Tversky Index and Focal Loss:
  - Tversky Index introduces a **weighting mechanism** to balance false positives (FP) and false negatives (FN).
  - Focal Loss amplifies the penalty for misclassifications, particularly for the minority class (fire regions).

Results: Fire Mask Prediction (U-Net)

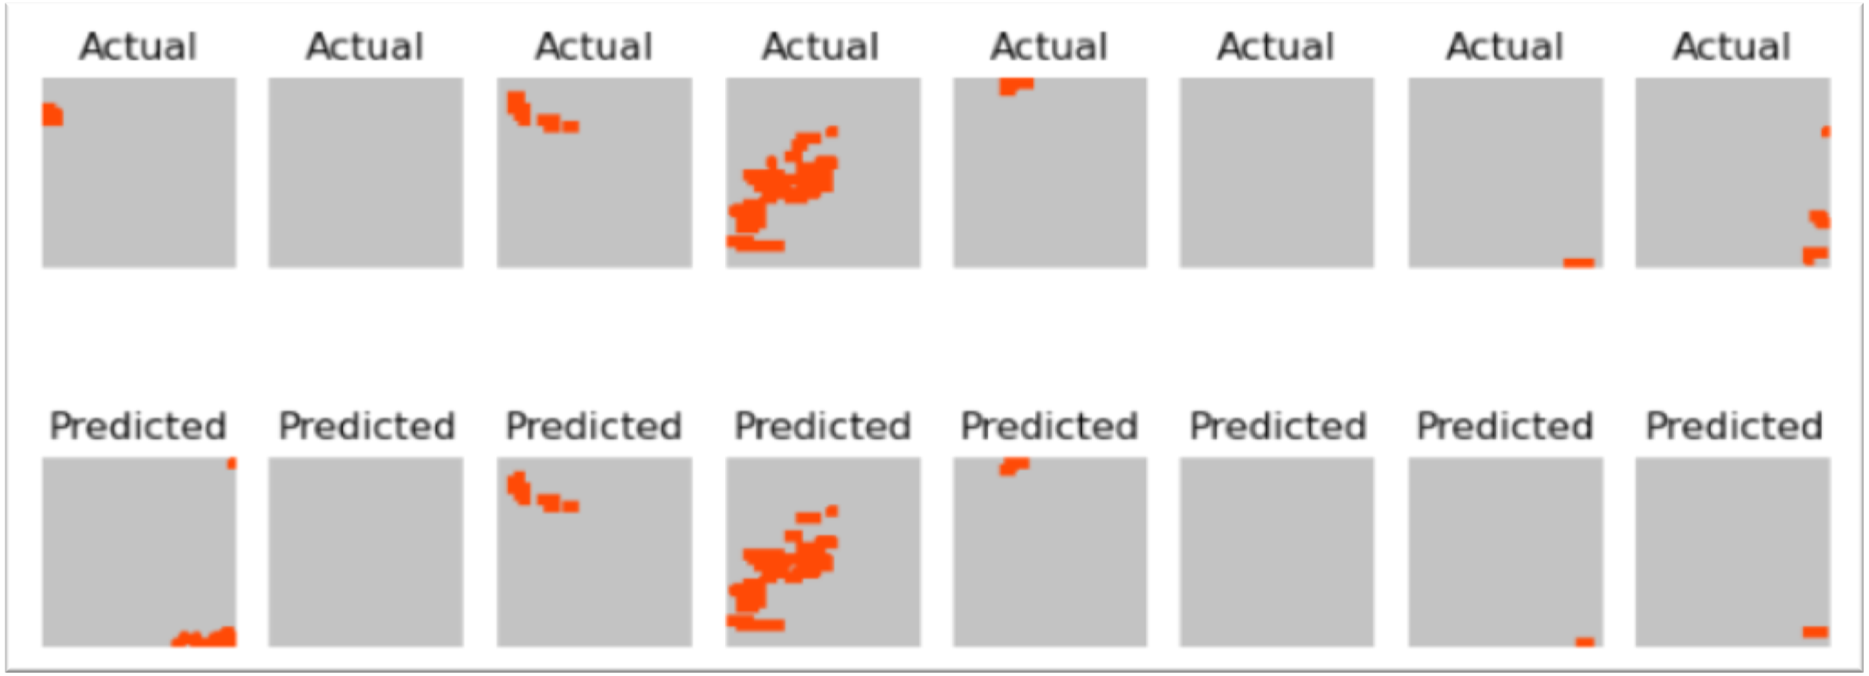

Predicted smaller fire regions well but struggled with spatial continuity in larger fire clusters, leading to gaps in prediction.

Results: Fire Mask Prediction (ViT)

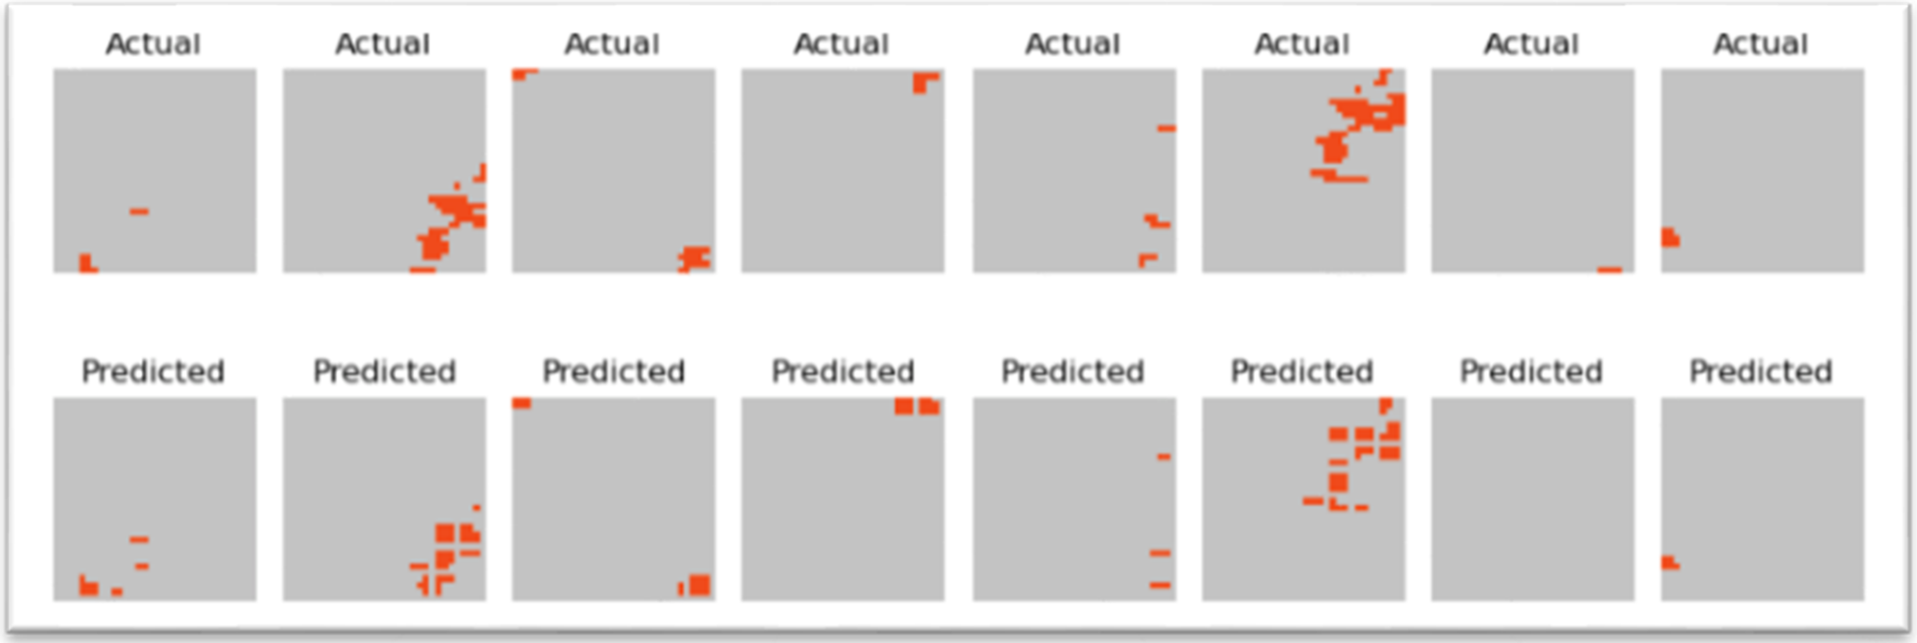

Captured larger fire clusters effectively but missed finer details, resulting in fragmented predictions for smaller fire regions.

Results: Fire Mask Prediction (Vi-Net)

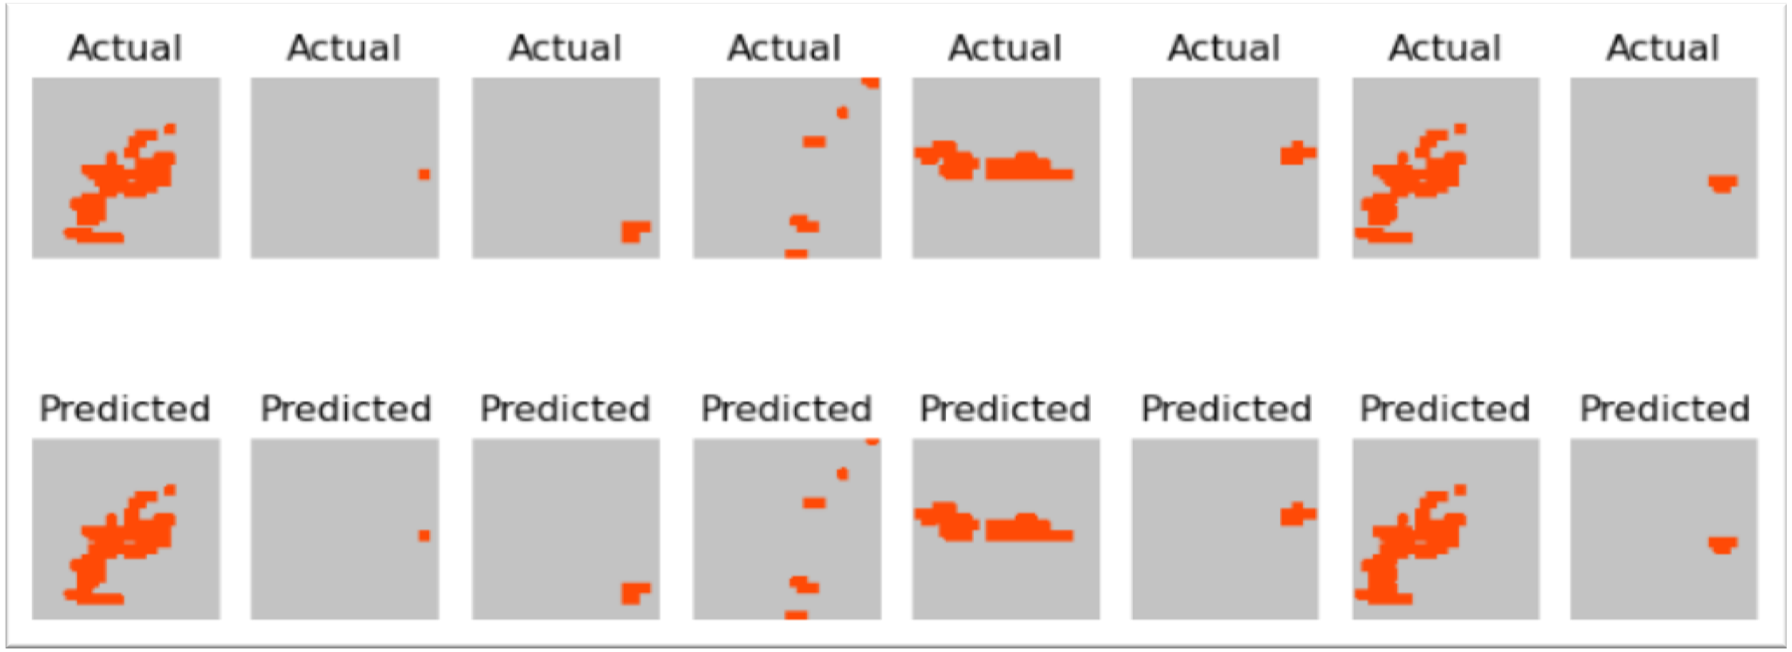

Achieved the most accurate predictions by integrating local precision (from U-Net) with global coherence (from ViT), reducing both under-prediction and over-prediction errors.

Results: Vi-Net Boundary Analysis

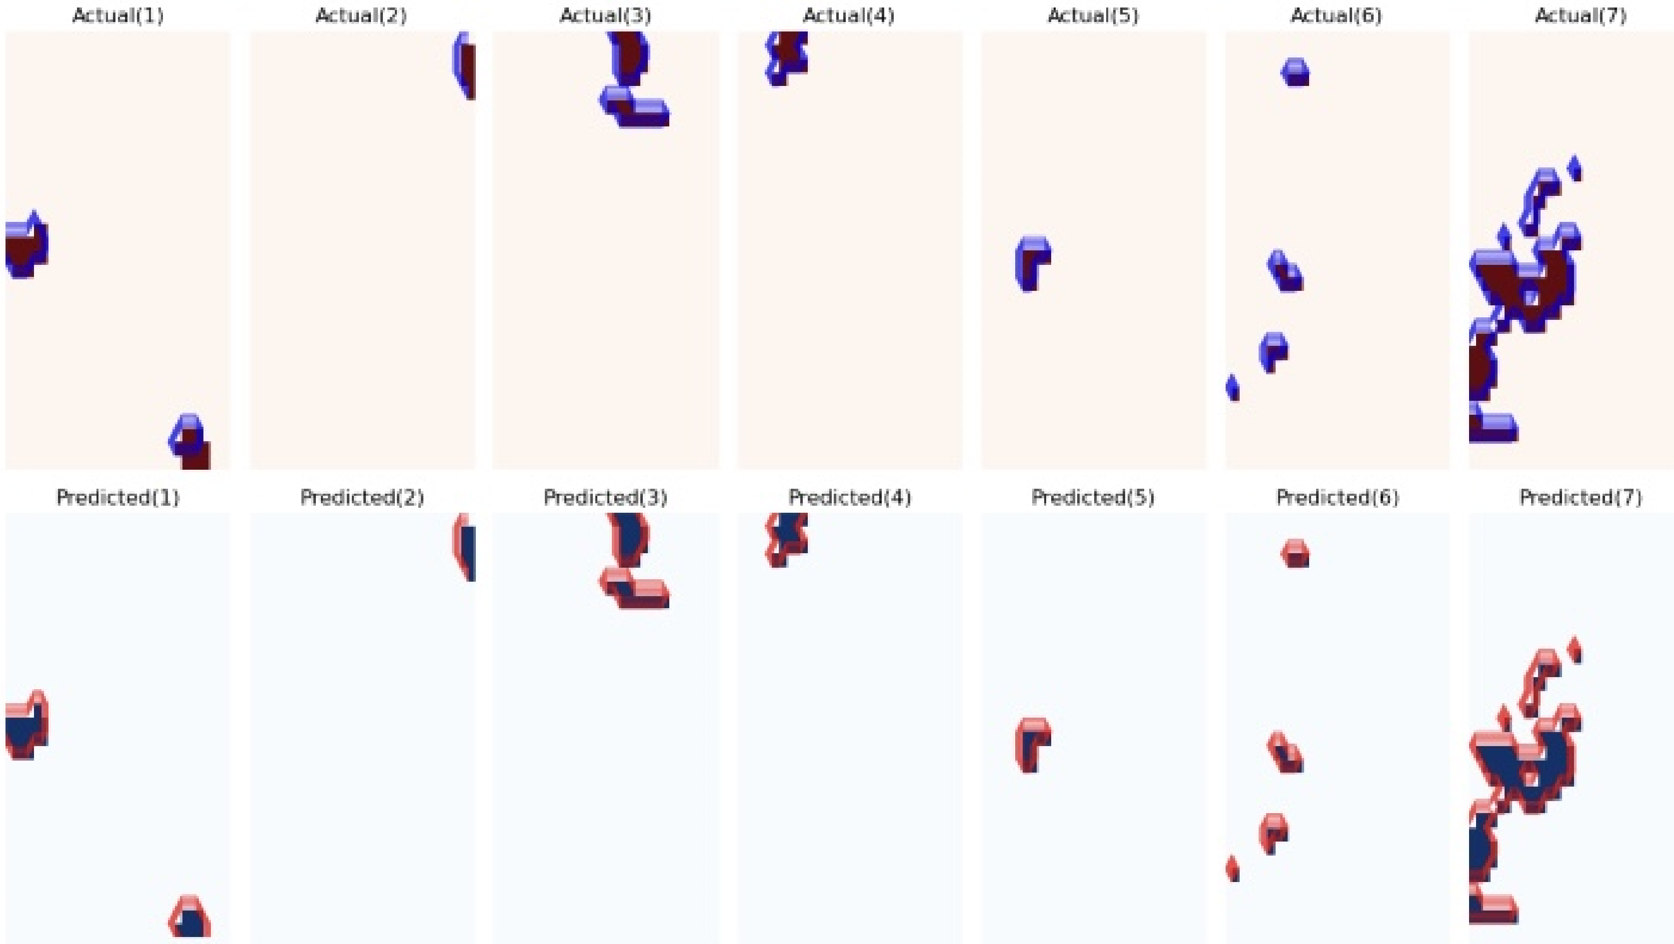

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Results: Performance Metrics

Dataset	U-Net				ViT				Vi-Net			
	IoU	F1	P.	R.	IoU	F1	P.	R.	IoU	F1	P.	R.
Training set	0.7879	0.8813	0.9981	0.7890	0.5625	0.7200	0.9130	0.5943	0.9803	0.9900	0.9979	0.9822
Validation set	0.7409	0.8511	0.9727	0.7566	0.5261	0.6894	0.9527	0.5402	0.9902	0.9880	0.9951	0.9809
Test set	0.7325	0.8403	0.9542	0.7506	0.5829	0.7418	0.9873	0.5941	0.9415	0.9725	0.9834	0.9619

- Vi-Net consistently outperforms both base models across all metrics with particularly strong F1 scores >0.97 on all datasets
- Notable improvement in Intersection over Union (IoU) from U-Net (0.73) and ViT (0.74) to Vi-Net (0.98) on test set, showing significantly better spatial accuracy
- Vi-Net achieves balanced Precision-Recall trade-off (both >0.96)

Results: Loss and Accuracy Graphs

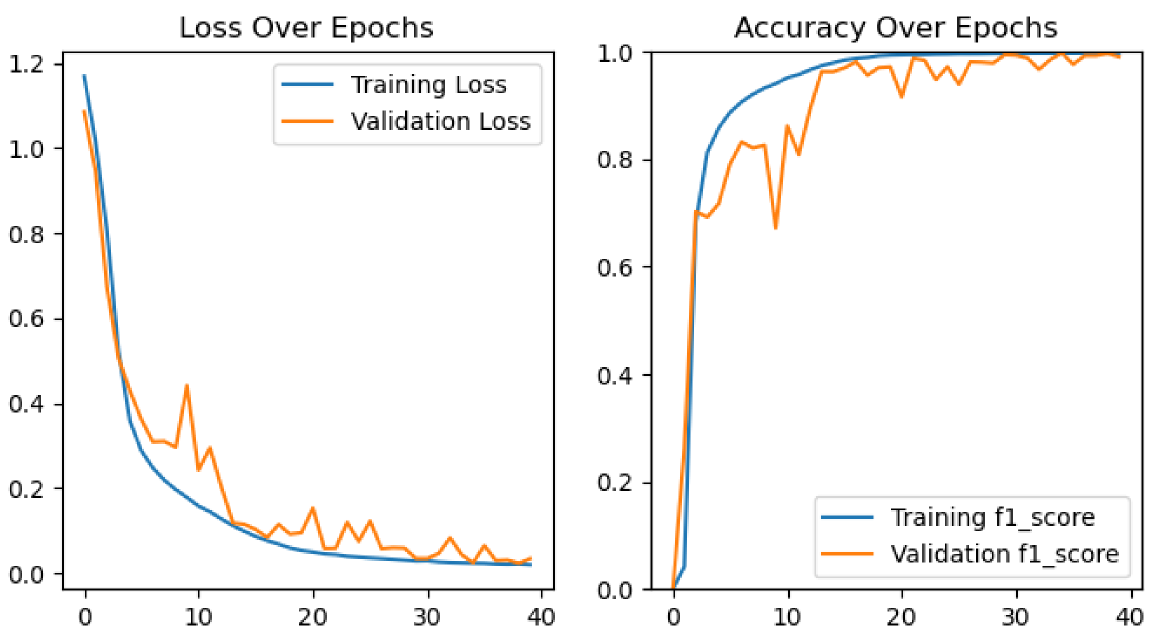

- Sharp loss reduction (0-10 epochs) → Fast learning
- Validation loss stability after epoch 20
- F1 score > 0.95 maintained
- Small train-val gap: Good generalization
- Consistent high F1 despite val loss fluctuations

Results: OA* Algorithm Implementation on Binary Grid

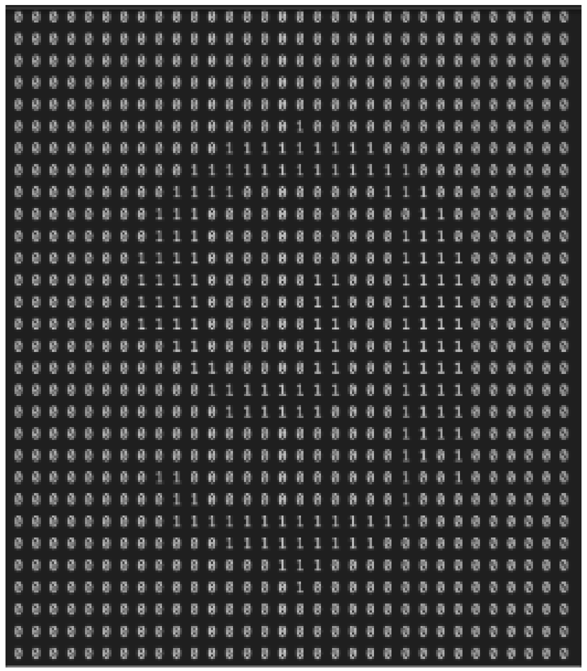

Binary Grid

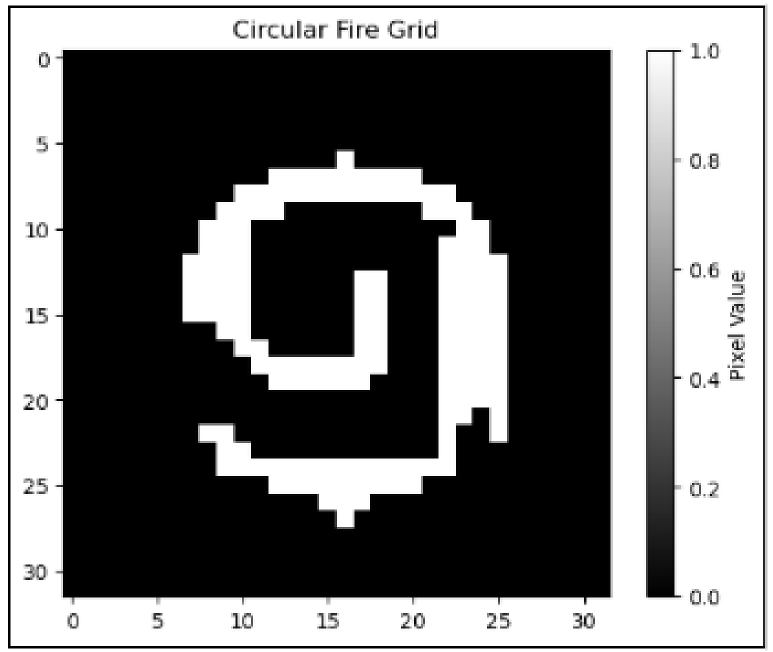

Corresponding Fire Mask

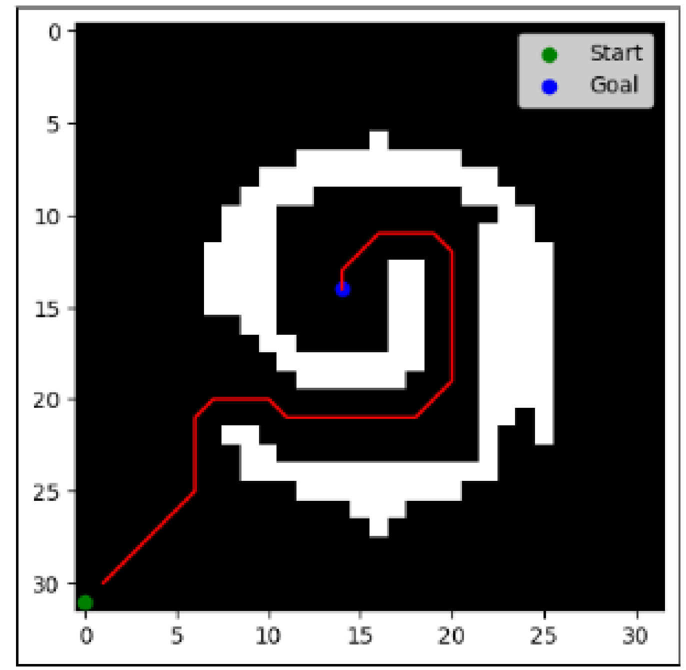

Generated Safe Path

Results: Scenarios with Navigable Safe Routes

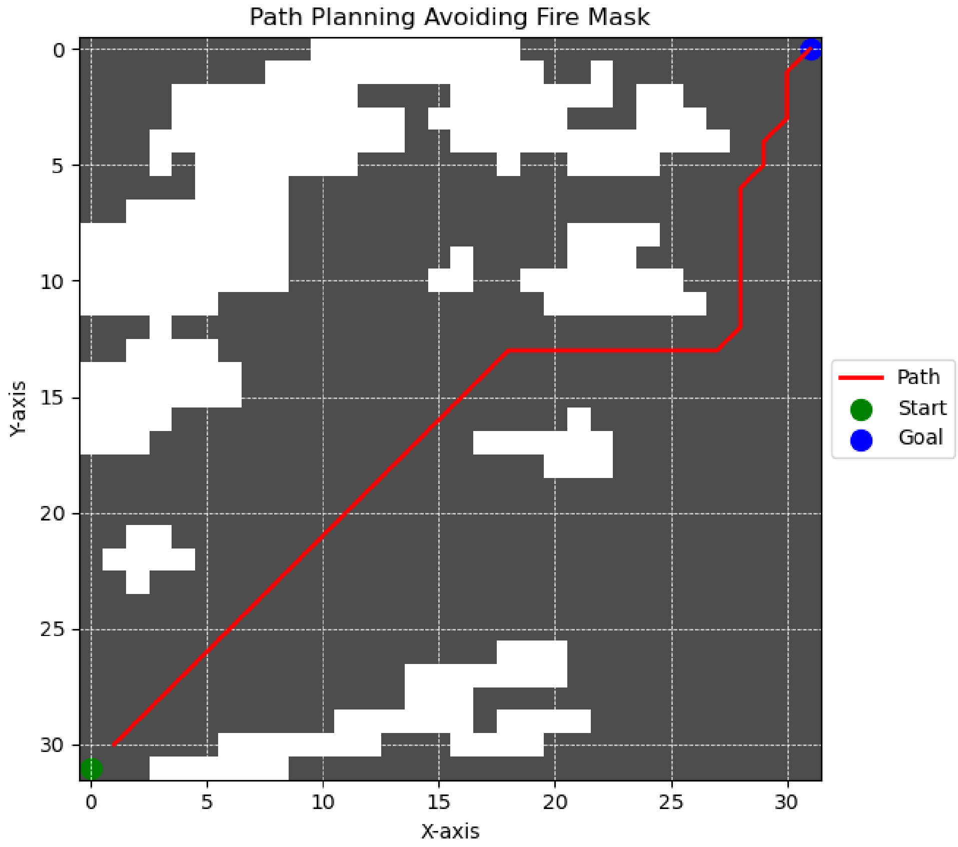

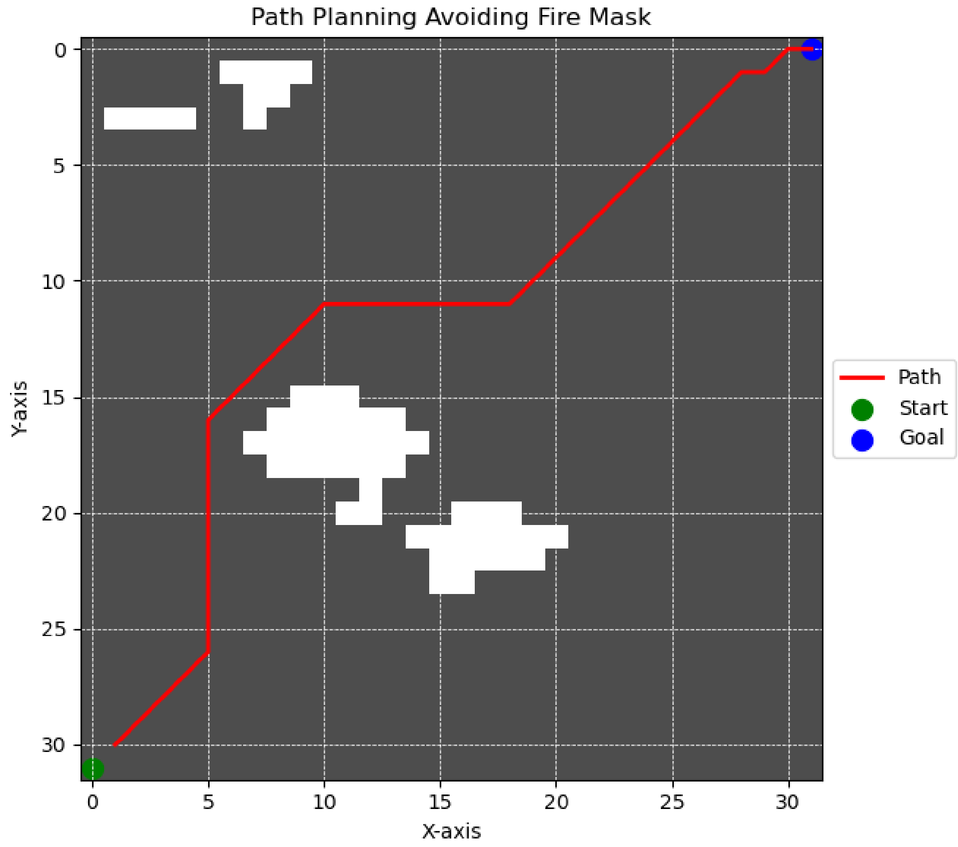

Results: Scenarios with Navigable Safe Routes

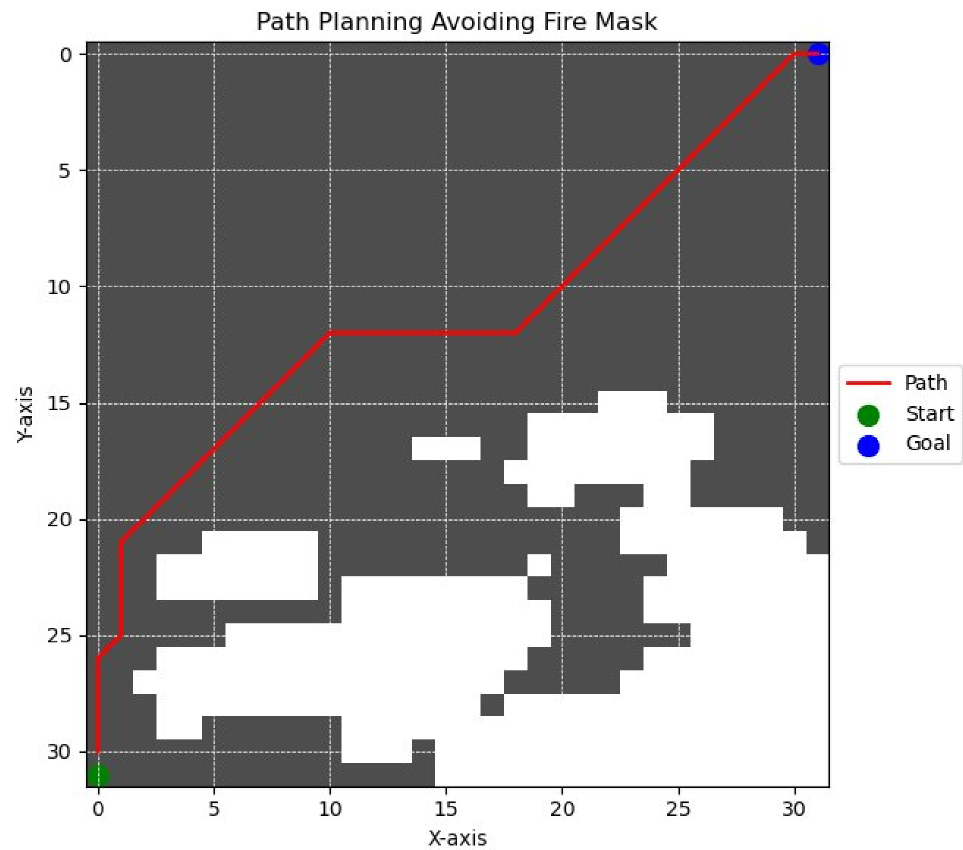

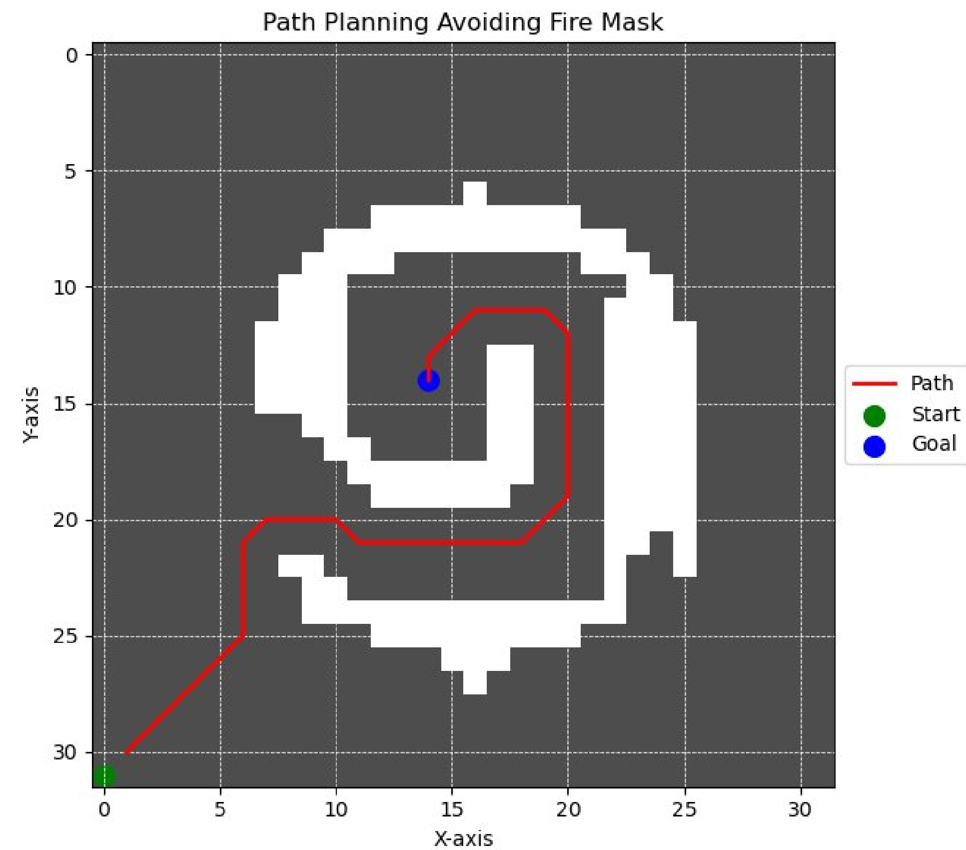

Results: Scenarios with No Safe Routes

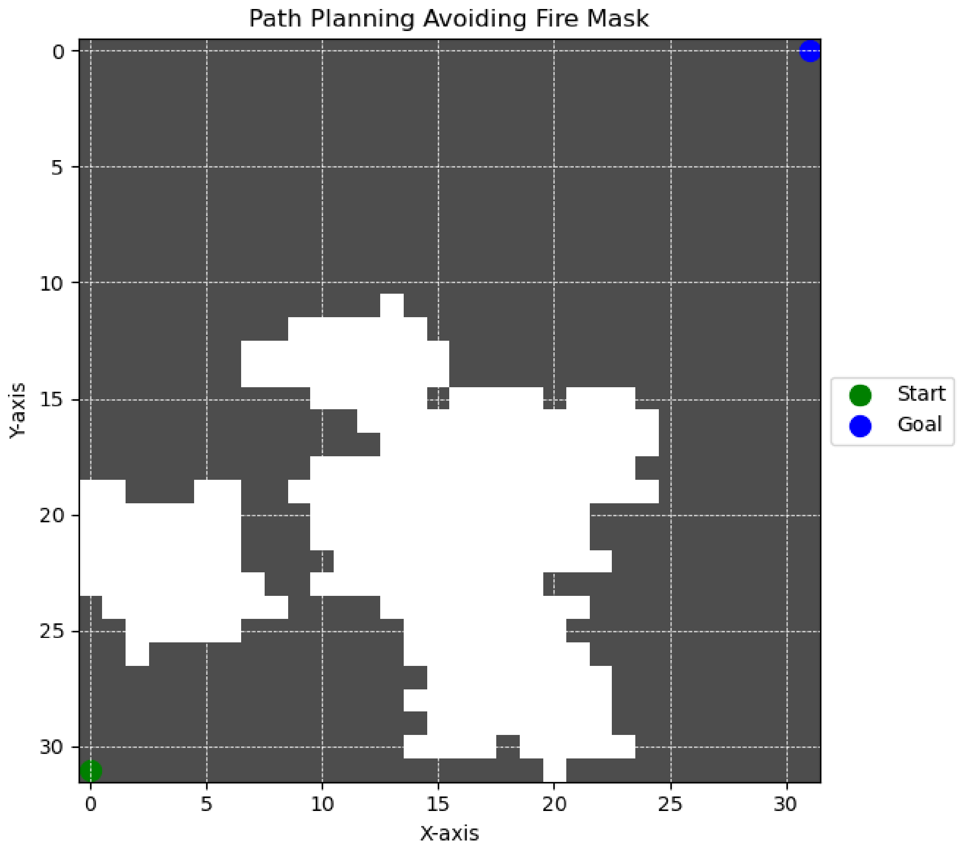

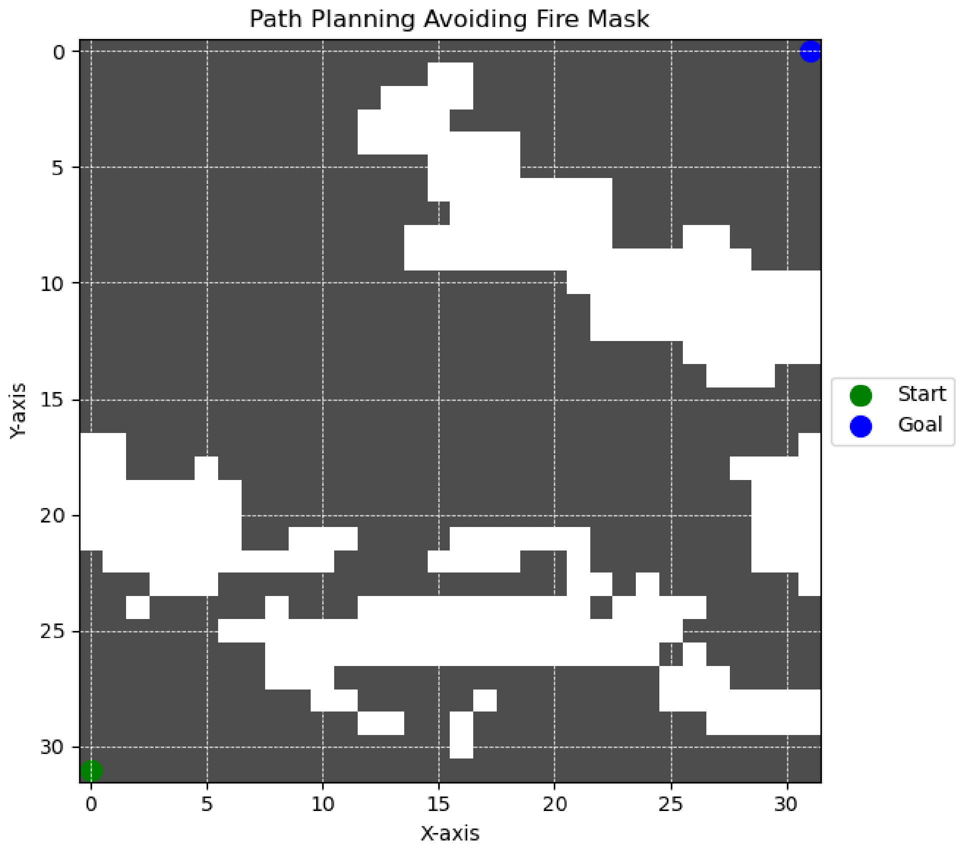

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Results: Scenarios with No Safe Routes

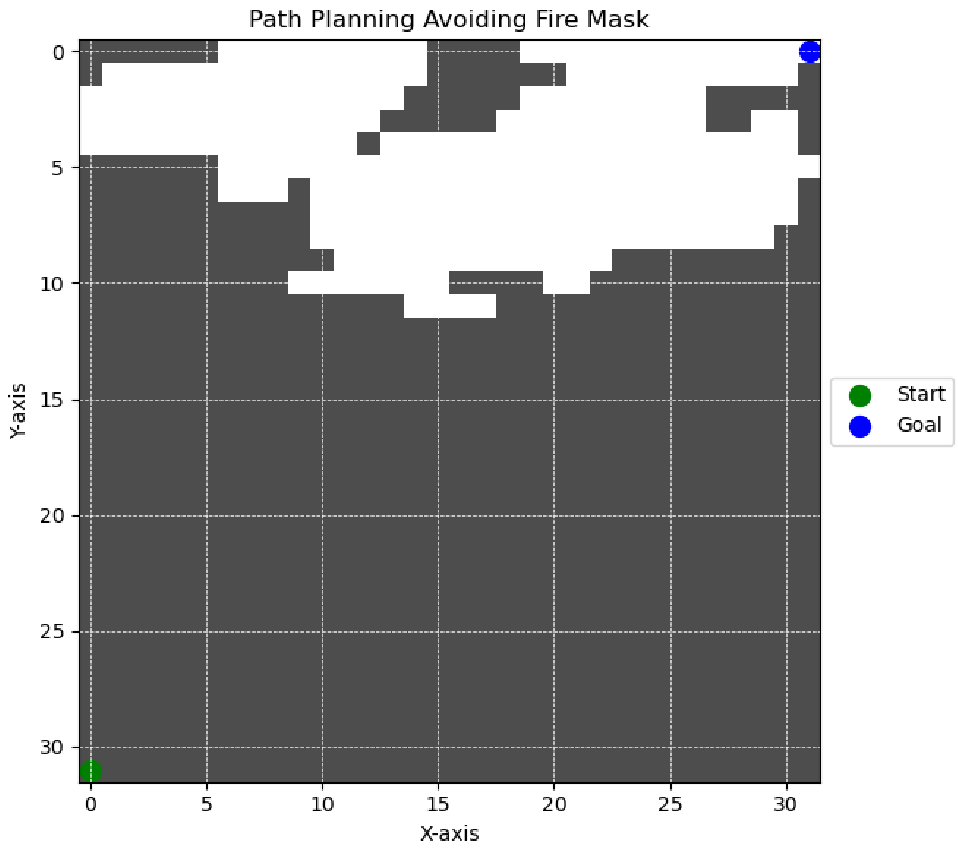

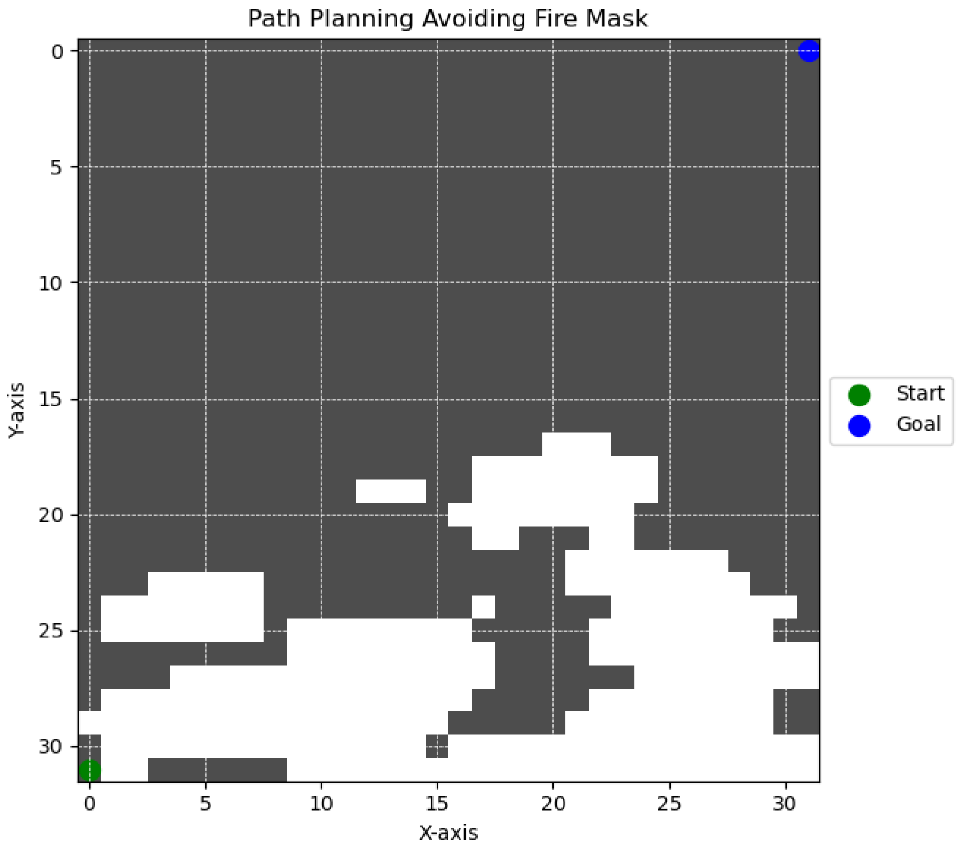

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

Results: Other Model Learnings

GCN Model

- An attributed graph is a static graph that associates each node with a set of attributes, representing node features.
- Each pixel = node in graph
- Define a correlation threshold. If the correlation between two nodes exceeds this threshold, an edge is created between them.

Performance	Metrics
Predictions	Not accurate
F1 Score	51.20%

MPL gather current information of neighbor nodes, combine it to get a new embedding, and update node embeddings.

Aggregate:

For each node, the model gathers feature information from its neighboring nodes (those connected by edges).

Then, it **aggregate** the features from the neighbors. This could be a simple sum, mean, or a learned weighted sum of the neighbor features.

Update:

Each node **updates** its feature vector based on the aggregated information and its own features.

This step involves passing the aggregated features through neural network layers (e.g., GCNConv layers).

Results: Other Model Learnings

Spatial Context Preservation:

- GNNs, while powerful in capturing relational data, might struggle to preserve local spatial contexts effectively when image data is transformed into graphs. Unlike CNNs or ViTs, GNNs do not inherently understand Euclidean space (common in image data), making them less intuitive for tasks requiring awareness of spatial organization directly from raw data formats like images or grids.

Graph Construction Limitations:

- The effectiveness of GNNs heavily depends on the quality of the graph construction. Other methods for graph creation: region-based, feature clustering. There must be check employed to check if the graph exhibits expected patterns (e.g., nodes representing similar image regions should have higher connectivity). Check: degree distribution, clustering coefficients

Conclusion and Future Work

CONCLUSION

Hybrid Vi-Net Model: Combines U-Net and ViT strengths, achieving 97.25% F1-score and 94.15% IoU, improving wildfire spread prediction.

Effective Loss Function: Focal Tversky Loss enhances model performance by addressing class imbalance and focusing on critical fire boundaries.

Safe Path Planning: Integrated predictive models with A* algorithm to create safe evacuation routes, ensuring reliable disaster mitigation.

Innovative Approach: Advanced machine learning methodologies set a new benchmark in predictive accuracy and real-world application for wildfire management.

Future Directions: Potential for real-time data integration to improve scalability and adaptability in dynamic wildfire scenarios.

Future Work

Extending the Dataset

Prediction Algorithm Approaches

Path Planning Algorithm Approaches

Integration of Real-Time Data Streams

Emergency Response Systems Enhancement

Interdisciplinary Collaboration

Explainability and Interpretability

THANK YOU

UNIVERSITY OF WATERLOO | FACULTY OF ENGINEERING

QUESTIONS & DISCUSSION

UNIVERSITY OF WATERLOO

FACULTY OF ENGINEERING

YOU+WATERLOO

Our greatest impact happens together.