Deep Transfer Learning for Early Wildfire Detection Using Multi-Spectral Satellite Imagery

Sarah Chen, Ph.D.
Nexa Research Institute
Machine Learning & Remote Sensing Division

James Hartley, M.Sc.
Nexa Research Institute
Edge Computing Systems Lab

Maria Kowalski, Ph.D.
Meridian University
Department of Environmental Data Science

ISCES 2026
International Symposium on Computational Environmental Science • Barcelona, Spain

94.7% Accuracy

3.2 min Latency

50K Images

Keywords: Wildfire detection, transfer learning, satellite imagery, edge computing, attention mechanisms

Abstract & Background

Early detection of wildfires is critical for minimizing environmental damage and protecting vulnerable communities. Traditional detection methods relying on human observation or ground-based sensor networks suffer from limited coverage and significant response delays. This study presents a novel deep transfer learning framework that leverages multi-spectral satellite imagery for real-time wildfire detection at a global scale.

Our approach combines pre-trained convolutional neural networks with spatial-channel attention mechanisms and domain-specific fine-tuning on a curated dataset of over 50,000 satellite images spanning five continents and seven distinct climate zones. The framework achieves 94.7% detection accuracy with a mean detection latency of 3.2 minutes when deployed on edge computing hardware.

Motivation

Climate change has intensified wildfire frequency and severity worldwide. In 2024 alone, wildfires burned over 12 million hectares globally, emitting an estimated 1.8 gigatons of CO_2. Current satellite-based detection systems, such as MODIS and VIIRS, provide valuable data but suffer from latency issues and limited spatial resolution, often detecting fires only after they have grown to substantial sizes.

Research Gap

While deep learning has shown promise in remote sensing applications, existing models are typically trained on region-specific data and fail to generalize across diverse ecosystems. Furthermore, most approaches require cloud-based inference pipelines, introducing unacceptable latency for time-critical wildfire response scenarios.

Key Challenge: Can we build a lightweight, globally-generalizable deep learning model that detects wildfires within minutes using standard satellite imagery, while running on edge devices at ground stations?

Objectives & Contributions

Research Objectives

O1 Develop a lightweight deep learning architecture optimized for real-time wildfire detection from multi-spectral satellite imagery.

O2 Leverage transfer learning to enable rapid model adaptation across geographically and climatically diverse regions with minimal labeled data.

O3 Optimize the inference pipeline for deployment on edge computing devices co-located with satellite ground stations.

O4 Comprehensively validate detection performance against existing operational systems including MODIS and VIIRS active fire products.

Key Contributions

C1. A novel attention-enhanced transfer learning architecture that improves cross-region generalization by 18.4% over standard fine-tuning.

C2. The WildFire-Net dataset: 50,000 curated multi-spectral image patches with pixel-level fire masks across five continents.

C3. An edge-optimized inference pipeline achieving 3.2-minute end-to-end latency on Lumina Jetson Orin hardware, enabling on-site deployment.

C4. Comprehensive benchmark across seven climate zones demonstrating consistent performance (F1 > 0.89 in all regions) and a 12.3% improvement over the best existing CNN-based method.

Research Questions

RQ1. How effectively can transfer learning bridge the domain gap between pre-trained vision models and multi-spectral satellite data?

RQ2. What is the minimum spatial and temporal resolution required for reliable early-stage wildfire detection?

RQ3. Can model compression techniques reduce inference latency below 5 minutes without sacrificing detection accuracy?

Methodology

Dataset: WildFire-Net

We curated a dataset of 50,000 multi-spectral image patches (512×512 pixels) from Sentinel-2 and Landsat-8 satellites. Images span 2019–2025 and cover boreal forests (Canada, Siberia), Mediterranean shrublands (Spain, Greece), tropical savannas (Brazil, Australia), temperate forests (California, Portugal), and arid grasslands (Central Asia). Each patch is labeled with a binary fire/non-fire mask verified by expert annotators.

Model Architecture

Our framework, named **WildFire-Net**, builds on EfficientNet-B4 pre-trained on ImageNet-21K. We augment the backbone with:

- **Spatial Attention Module:** 3×3 depthwise-separable convolutions with dilated kernels (rates = 1, 2, 3) to capture multi-scale smoke patterns.
- **Channel Attention Module:** Squeeze-and- excitation blocks recalibrating spectral band weights, emphasizing near-infrared (B8) and shortwave-infrared (B11, B12) bands critical for fire detection.
- **Lightweight Decoder:** A single upsampling block with skip connections, reducing parameters by 40% compared to U-Net-style decoders.

Training Protocol

Two-stage training on $4 \times$ Lumina A100 GPUs:

Stage 1 — Pre-training: ImageNet-21K weights fine-tuned on the BigEarthNet-MM multi-spectral benchmark (590K Sentinel patches, 43 land-cover classes) for 50 epochs. Learning rate: 1e-4 with cosine annealing.

Stage 2 — Domain Adaptation: Fine-tuned on WildFire-Net with a weighted focal loss ($\gamma = 2.0$, $\alpha = 0.75$ to address class imbalance). Data augmentation includes random rotations, flips, spectral jitter, and CutMix. Batch size 64, 30 epochs.

Evaluation Metrics

Precision, Recall, F1-score, Intersection-over-Union (IoU), and end-to-end detection latency measured from image acquisition to alert generation.

Results & Analysis

Detection Performance

Method	Prec.	Rec.	F1	IoU
MODIS C6	0.789	0.717	0.751	0.602
VIIRS AF	0.821	0.754	0.786	0.648
ResNet-50	0.864	0.828	0.845	0.732
U-Net++	0.891	0.863	0.877	0.781
WildFire-Net	**0.932**	**0.918**	**0.925**	**0.860**

Table 1. Performance comparison on the WildFire-Net test set (10,000 images).

Per-Region Generalization

WildFire-Net maintains strong performance across all seven climate zones (F1 range: 0.892–0.948), compared to ResNet-50 (F1 range: 0.731–0.864). The attention modules are particularly effective in Mediterranean and boreal regions where smoke dispersion patterns differ from temperate training data.

Key Finding: The spatial-channel attention mechanism improves cross-region F1 by 8.3 percentage points compared to the same architecture without attention, confirming that attention helps the model focus on fire-specific spectral signatures rather than region-specific background features.

Latency Analysis

Platform	Latency	Throughput
Cloud (A100)	0.82 s	1,220 img/h
Edge (Jetson)	3.2 min	18.8 img/h
CPU (Xeon)	8.7 min	6.9 img/h

Table 2. Inference latency across deployment platforms. Edge latency meets the 5-minute operational target for ground-station

Conclusions & Resources

Summary

This study demonstrates that deep transfer learning with attention mechanisms can achieve state-of-the-art wildfire detection from satellite imagery while maintaining sufficient efficiency for edge deployment. WildFire-Net outperforms existing operational systems by 12–18% and generalizes robustly across diverse ecosystems.

Limitations

- Detection of small fires (< 0.5 ha) remains challenging due to spatial resolution constraints (10–30 m per pixel).
- Nighttime detection relies exclusively on thermal bands, reducing accuracy to ∼82%.
- Cloud cover obscures optical imagery; integration with SAR data is needed for all-weather capability.

Future Work

Drone Integration: Extending the framework to ingest high-resolution UAV imagery for verification of satellite-based detections.

Multi-Modal Fusion: Combining optical, thermal, and synthetic aperture radar (SAR) data for robust all-weather, day-night detection.

Operational Pilot: Planned deployment at three ground stations in partnership with regional fire management agencies in California, Portugal, and Australia (2026–2027).

References

[1] S. Chen & J. Hartley. Transfer learning strategies for multi-spectral Earth observation. *IEEE Trans. Geosci. Remote Sens.*, 63:1–18, 2025.

[2] M. Kowalski, S. Chen, & J. Hartley. Attention mechanisms for environmental monitoring with satellite imagery. *NeurIPS 2024 Workshop on Climate Change AI*, Vancouver, 2024.

[3] Nexa Research Institute. WildFire-Net: An edge-AI framework for environmental threat detection. Technical Report NX-TR-2025-042, 2025.

[4] L. Zhang & K. Park. BigEarthNet-MM: A large-scale multi-modal benchmark for remote sensing. *CVPR 2024*, Seattle, 2024.

Acknowledgments

This work was supported by the Nexa Environmental AI Initiative (Grant NEAI-2025-042) and the Meridian University High-Performance Computing Center. We thank the European Space Agency and USGS for providing open-access satellite imagery.