INFERNIS

INTELLIGENCE FORGED IN FIRE

A Machine Learning Engine for Wildfire Prediction in British Columbia

White Paper — Version 1.0

Date	February 2026
Author	Arnold Dogelis
Contact	hello@argonbi.com
Location	Vancouver, BC

CONTENTS

Table of Contents

1. Executive Summary
2. The Problem: Wildfire Risk in British Columbia
3. The Solution: INFERNIS
4. The Science
5. Empirical Feature Importance
6. Research Validation
7. Data Foundation
8. System Architecture
9. Risk Scoring Methodology
10. Use Cases
11. Competitive Landscape
12. Roadmap
13. Technical Specifications
14. Limitations and Uncertainties
15. Conclusion
16. References

SECTION 1

Executive Summary

INFERNIS is a machine learning engine purpose-built to predict wildfire occurrence and spatial risk across British Columbia. By combining gradient-boosted decision trees with convolutional neural networks in a regionally calibrated ensemble, INFERNIS ingests 21 open data sources – spanning satellite imagery, reanalysis weather grids, multi-depth soil moisture profiles, vegetation indices, and historical fire records – to produce calibrated daily fire occurrence probabilities at 1km grid resolution across the entire province.

The problem INFERNIS addresses is both urgent and structurally underserved. British Columbia has experienced three of its worst wildfire seasons in recorded history within the last decade, with cumulative damages exceeding tens of billions of dollars and hundreds of thousands of residents displaced. Existing fire danger systems, while scientifically rigorous, remain station-based and reactive. They were not designed to deliver the granular, probabilistic, machine-learning-enhanced predictions that modern fire management, insurance underwriting, and infrastructure planning demand. INFERNIS closes that gap: it transforms Canada's world-class open data ecosystem into actionable, API-delivered wildfire intelligence.

Key differentiators include BC-specific model calibration across 14 biogeoclimatic zones, a hybrid ensemble architecture validated through walk-forward temporal backtesting, automated daily ingestion of open data sources requiring zero manual intervention, and a REST API designed for integration by government agencies, insurers, utilities, and application developers. INFERNIS is built entirely on freely available Canadian and international open data.

SECTION 2

The Problem: Wildfire Risk in British Columbia

British Columbia is in a wildfire crisis that is accelerating, not stabilizing. The 2017 fire season burned 1.2 million hectares and was declared the worst in provincial history. That record fell the very next year when 2018 saw 1.35 million hectares consumed. Both were eclipsed by the 2023 season, which burned 2.84 million hectares – more than double the previous record and the worst wildfire season in Canadian history. This is not cyclical variation. It is a structural shift driven by climate change, fuel accumulation from decades of fire suppression, and expanding human-wildland interface.

The economic toll is staggering. The 2023 season is expected to have caused multi-billion-dollar losses when accounting for direct suppression expenditures (over $1 billion CAD in BC alone), insured property losses, infrastructure damage, public health impacts from smoke exposure, and economic disruption from evacuations that displaced over 35,000 residents. Insurance losses from wildfire are now a material line item in Canadian reinsurance portfolios, and multiple insurers have begun restricting coverage in high-risk BC communities.

Current wildfire danger assessment in Canada relies on the Canadian Forest Fire Danger Rating System (CFFDRS), a scientifically rigorous framework developed over decades by the Canadian Forest Service. The CFFDRS and its Fire Weather Index (FWI) System remain the gold standard for fire danger rating worldwide. However, the system was designed for an era of manual weather station readings and human interpretation. It is historically station-based, though gridded FWI products based on ERA5-style inputs now exist. The core operational system does not incorporate machine learning, satellite-derived vegetation indices, or multi-depth soil moisture reanalysis. The Canadian Wildland Fire Information System (CWFIS) does publish daily and forecast fire danger maps and exposes data layers via OGC web services (WMS/WFS/WCS), but it does not provide a modern JSON-based developer API. Its documentation and interfaces are oriented toward government fire management workflows and GIS professionals rather than commercial developers, insurers, or application builders.

The gap between what the science makes possible and what the operational systems deliver represents both a public safety risk and a commercial opportunity.

SECTION 3

The Solution: INFERNIS

INFERNIS is a machine learning-powered fire prediction engine designed specifically for British Columbia. It operates as a daily batch prediction system, ingesting data each afternoon after noon weather observations are finalized, and producing grid-level fire occurrence probabilities and danger classifications that are served via a REST API.

The architecture employs a hybrid ensemble approach. An XGBoost gradient-boosted classifier serves as the primary occurrence prediction model, trained on approximately 298,000 labeled samples (27,146 positive, 271,460 negative at a 10:1 ratio with spatiotemporal buffering) spanning eleven fire seasons (2015–2025) and incorporating 28 engineered features drawn from weather reanalysis, FWI components, multi-depth soil moisture, vegetation indices, topographic derivatives, infrastructure proximity, and lightning detection. In cross-validated evaluation on 1km grid data, the XGBoost model achieves an AUC-ROC of 0.974 – meaning it correctly ranks fire-prone conditions above non-fire conditions 97.4% of the time. A FireUNet convolutional neural network (7.8M parameters) operates in parallel, processing 12-channel spatial inputs across a 256×512 raster grid to generate continuous risk heatmaps that capture landscape-scale fire spread patterns and spatial autocorrelation that point-based models miss, achieving an AUC-ROC of 0.815. The Risk Fuser combines both model outputs in logit space with per-zone calibration coefficients independently tuned for each of BC’s 14 biogeoclimatic (BEC) zones, producing a final composite risk score for each grid cell.

Beyond current-day predictions, INFERNIS produces multi-day fire risk forecasts by ingesting numerical weather prediction (NWP) data from Environment and Climate Change Canada. High-Resolution Deterministic Prediction System (HRDPS) forecasts at 2.5km resolution drive the first two forecast days, while Global Deterministic Prediction System (GDPS) data extends predictions out to 10 days. FWI moisture codes are rolled forward day-by-day using forecast weather, maintaining physical consistency across the forecast horizon. A confidence decay factor attenuates predictions at longer lead times to transparently communicate increasing uncertainty.

All data sources are open and freely available from Canadian federal agencies, provincial data portals, and international scientific repositories. INFERNIS requires no proprietary data subscriptions. Google Earth Engine is used under its non-commercial license for satellite data access during development, with a plan to transition to a commercial license once revenue is generated. The system is designed for full automation: once configured, the daily pipeline runs without human intervention, from data retrieval through model inference to API delivery.

SECTION 4

The Science

The Fire Weather Index System

INFERNIS builds on the scientific foundation of the Canadian Fire Weather Index (FWI) System, which models fire danger through a three-tier structure of moisture codes and fire behavior indices.

At the base tier, three moisture codes track fuel drying at different time scales. The Fine Fuel Moisture Code (FFMC) represents the moisture content of surface litter and fine fuels, responding to weather changes within hours. The Duff Moisture Code (DMC) tracks moisture in loosely compacted organic layers, operating on a time scale of days to weeks. The Drought Code (DC) models deep organic soil moisture with a seasonal memory spanning weeks to months. Together, these three codes encode cumulative drying across the full spectrum of fuel layers relevant to fire ignition and behavior.

The intermediate tier combines these moisture codes into two compound indices. The Initial Spread Index (ISI) merges FFMC with wind speed to estimate the expected rate of fire spread. The Buildup Index (BUI) combines DMC and DC to represent the total fuel available for combustion.

At the top tier, the Fire Weather Index (FWI) integrates ISI and BUI into a single numeric rating of fire intensity. INFERNIS uses all six FWI components as engineered features in its ML models, preserving the decades of fire science encoded in their formulations while allowing the machine learning layer to discover nonlinear interactions and threshold effects that the linear FWI aggregation cannot capture.

Machine Learning Enhancement

Beyond the classical FWI features, INFERNIS incorporates modern data sources that were unavailable when the CFFDRS was designed. Satellite-derived vegetation indices (NDVI, EVI, LAI) from MODIS and Sentinel-2 provide direct observation of fuel condition, canopy structure, and vegetation stress. ERA5 reanalysis provides gridded, gap-free precipitation and evapotranspiration measurements alongside soil moisture at four depths spanning from the surface through to deep soil layers. High-resolution topographic features derived from the Canadian Digital Elevation Model (CDEM) – elevation, slope, aspect, and hillshade – encode terrain characteristics that influence fire behavior at the landscape scale. Distance to the nearest road, derived from the BC Digital Road Atlas, captures human-wildland interface proximity, a key factor in both ignition likelihood and suppression access. Lightning detection from the Canadian Lightning Detection Network (CLDN) captures the primary natural ignition source – lightning is responsible for roughly 60% of BC wildfire ignitions and, as in the rest of Canada, for the majority of the total area burned.

The complete 28-feature vector per grid cell per day comprises: 6 FWI components (FFMC, DMC, DC, ISI, BUI, FWI), 10 weather variables (temperature, relative humidity, wind speed, wind direction, 24h precipitation, soil moisture at four depths, evapotranspiration), 3 vegetation indices (NDVI, snow cover fraction, leaf area index), 5 topographic and infrastructure features (elevation, slope, aspect, hillshade, distance to nearest road), 2 temporal encodings (day-of-year sine and cosine), and 2 lightning features (24h and 72h flash density).

Training Data Construction

Labels are derived from the Canadian National Fire Database (CNFDB) point-of-origin records and BC Wildfire Service perimeter data. A grid cell is labeled positive for a given day if a fire ignition point falls within its 1km boundary during that day, or if the cell is within a 3km spatial buffer and 3-day temporal window of an ignition point. This buffering reflects the reality that fires do not ignite instantaneously at a single point – ignition conditions exist across a spatial neighborhood before and during the event. Negative samples are drawn from fire-free cells with spatiotemporal exclusion: negatives must be at least 10km and 7 days from any fire event, preventing contamination from near-miss conditions that are functionally fire-prone. The negative-to-positive ratio is 10:1 (271,460 negatives to 27,146 positives in the 1km training set), sampled with stratification across years and BEC zones to prevent temporal or geographic bias from dominating the training signal.

This label construction means that INFERNIS predicts fire occurrence conditions (whether a cell exhibits the combination of weather, fuel, and landscape factors associated with nearby ignition) rather than predicting the precise location of a specific fire start to the exact grid cell. The spatiotemporal buffering also serves as a data leakage prevention mechanism: by excluding near-fire cells from the negative pool, the model cannot learn trivially from spatial autocorrelation of adjacent fire/non-fire cells.

SECTION 5

Empirical Feature Importance

Training on the 1km grid across eleven BC fire seasons with the full 28-feature vector reveals the following feature importance ranking by mean |SHAP| value:

RANK	FEATURE	MEAN SHAP	CATEGORY
1	NDVI (vegetation greenness)	1.25	Vegetation
2	Elevation	1.03	Topography
3	DMC (duff moisture code)	0.88	FWI
4	DC (drought code)	0.74	FWI
5	Soil moisture	0.49	Weather
6	Day-of-year	0.44	Temporal
7	FFMC (fine fuel moisture code)	0.42	FWI
8	Temperature	0.31	Weather
9	Wind speed	0.28	Weather

Vegetation condition (NDVI) remains the single most important predictor. Elevation ranks second, demonstrating that topographic context materially improves prediction at 1km resolution. Three FWI components (DMC, DC, FFMC) appear in the top nine, confirming the value of the classical fire weather indices as ML features. Soil moisture ranks fifth, capturing landscape-level dryness that integrates weeks of precipitation and evapotranspiration history.

SECTION 6

Research Validation

The approach is validated by a growing body of peer-reviewed research. Recent machine learning studies on regional wildfire prediction using gradient-boosted models with ERA5 and FWI feature sets report AUCs in the 0.8–0.9 range, depending on region, task definition, and evaluation protocol. The BCWildfire benchmark dataset, which evaluates deep learning models at 1km resolution across British Columbia, reports that recent architectures (e.g., S-Mamba) achieve F1 scores above 0.85 and PR-AUC close to 0.95 in boreal wildfire risk prediction.

INFERNIS achieves an AUC-ROC of 0.974 and average precision of 0.794 in 5-fold stratified cross-validation on its 1km, 11-year training corpus, with a Brier score of 0.036 indicating well-calibrated probability outputs. Walk-forward temporal backtesting (training on years [2015, test_year-1], testing on test_year) yields AUC-ROC values of 0.90–0.93 across six held-out fire seasons (2019–2024), confirming that model performance generalizes across years and is not an artifact of random cross-validation splits.

SECTION 7

Data Foundation

INFERNIS draws from 21 open data sources spanning eight major categories. A comprehensive catalog is maintained in the project's DATA_SOURCES.md document; the following summarizes the key inputs.

Historical Fires

The Canadian National Fire Database (CNFDB) provides point-of-origin records for fires dating back decades. BC Wildfire Service perimeter data supplies polygon boundaries for all significant fires, enabling both point-based classification training and spatial burn-area modeling.

Weather

ERA5 reanalysis from the European Centre for Medium-Range Weather Forecasts (ECMWF) serves as the primary weather backbone. ERA5 provides hourly, gridded, gap-free atmospheric variables at approximately 31km native resolution, globally, from 1940 to present with a five-day lag. INFERNIS ingests 2m temperature, dewpoint, 10m wind components, total precipitation, potential evapotranspiration, and soil moisture at multiple depths. Precipitation and evapotranspiration rank among the top 10 most important predictive features, confirming the value of ERA5 variables beyond what the classical FWI system incorporates.

Satellite Imagery

MODIS and Sentinel-2 imagery, accessed via Google Earth Engine, provides vegetation indices (NDVI, EVI), active fire detections (MODIS Thermal Anomalies), and burn severity assessments.

Soil Moisture

ERA5-Land soil moisture layers at four depth levels (0–7cm, 7–28cm, 28–100cm, 100–289cm) provide gridded subsurface water content critical for predicting deep organic fuel drying. All four layers are ingested as model features, capturing moisture gradients from surface litter through to deep root zones.

Vegetation and Fuel

NDVI, EVI, and Leaf Area Index (LAI) time series characterize vegetation phenology, canopy density, and stress. CFFDRS Fuel Behaviour Prediction (FBP) system fuel type maps classify the landscape into standardized fuel categories.

Topography

The Canadian Digital Elevation Model (CDEM) at approximately 23m resolution provides elevation, from which slope gradient, aspect angle, and hillshade illumination index are derived via numerical gradient computation. Elevation ranks as the 2nd most important feature in the trained 1km model (mean |SHAP| = 1.03), demonstrating that topographic context is a critical contributor to fire prediction.

Infrastructure

The BC Digital Road Atlas provides road network geometry used to compute distance-to-nearest-road for each grid cell, capturing human-wildland interface proximity relevant to both ignition probability and suppression accessibility.

Lightning

The Canadian Lightning Detection Network (CLDN) provides lightning strike locations and polarity, as lightning is responsible for roughly 60% of BC wildfire ignitions and, as in the rest of Canada, for the majority of the total area burned.

All data is sourced from Canadian federal agencies, the Government of British Columbia, ECMWF, NASA, and ESA. No proprietary data subscriptions are required.

SECTION 8

System Architecture

INFERNIS is organized into six core subsystems that execute as a coordinated daily pipeline.

DATA FORGE

DATA FORGE is the automated ingestion layer. It retrieves, validates, and standardizes data from all sources on a daily schedule, handling format conversions, coordinate reprojection, temporal alignment, and quality control. Data Forge maintains a local mirror of key datasets and performs incremental updates to minimize bandwidth and processing time.

FIRE CORE

FIRE CORE is the primary prediction engine, built on XGBoost. It operates on a structured 28-feature matrix with one row per grid cell per day, incorporating FWI components, weather variables, multi-depth soil moisture, vegetation indices, topographic features, infrastructure proximity, lightning activity, and temporal encodings. Trained on approximately 298,000 labeled samples (10:1 negative-to-positive ratio with spatiotemporal buffering) spanning 2015–2025 at 1km resolution, the model achieves an AUC-ROC of 0.974 and a Brier score of 0.036 in cross-validated evaluation, producing well-calibrated occurrence probabilities.

HEATMAP ENGINE

HEATMAP ENGINE employs a FireUNet convolutional neural network (7.8M parameters) that processes 12-channel spatial inputs at 256×512 pixel resolution – stacked grids of weather, moisture, vegetation, and topography covering the full BC extent – to generate continuous spatial risk surfaces. The CNN captures landscape-scale patterns, spatial gradients, and neighborhood effects that the cell-independent XGBoost model cannot represent.

RISK FUSER

RISK FUSER combines outputs from FIRE CORE and HEATMAP ENGINE using a weighted ensemble operating in logit space with regional calibration. Calibration coefficients are fitted independently for each of BC's 14 biogeoclimatic (BEC) zones via logistic regression, accounting for dramatically different fire regimes, fuel types, and climate characteristics. The fuser transforms model outputs to logit space, applies zone-specific linear calibration, and maps the resulting probabilities to a six-level danger classification. In current calibration, XGBoost dominates the ensemble weighting across most zones – the CNN spatial risk component receives minimal weight in the logistic regression fit, indicating that the cell-level XGBoost predictions already capture most of the predictive signal at 1km resolution. The CNN architecture remains in the pipeline as an active component; future work on higher-resolution spatial inputs and more expressive CNN training may unlock additional ensemble value.

FORECAST ENGINE

FORECAST ENGINE extends predictions beyond the current day by combining high-resolution HRDPS weather forecasts (days 1–2, 2.5km resolution from Environment and Climate Change Canada) with global GDPS forecasts (days 3–10, 25km resolution). The engine rolls forward FWI moisture codes day-by-day using forecast weather, builds the full 28-feature matrix for each lead day, and applies the XGBoost model to produce multi-day fire risk trajectories. A confidence decay factor (default 0.95 per lead day) attenuates predictions at longer lead times to reflect increasing forecast uncertainty. Forecast weather is bilinearly interpolated from the native NWP grid to the INFERNIS prediction grid using scipy's RegularGridInterpolator.

REST API

REST API serves pre-computed daily predictions and multi-day forecasts via a FastAPI application. Consumers can query fire risk by geographic coordinates, grid cell identifiers, or bounding boxes, receiving JSON responses containing occurrence probabilities, danger classifications, contributing factor breakdowns, and up to 10 days of forecast trajectories with per-day confidence scores and data source attribution (HRDPS or GDPS).

The full pipeline executes daily at 14:00 Pacific Time, after noon weather observations have been incorporated into data feeds. The daily pipeline produces current-day predictions; the forecast pipeline runs immediately afterward, generating multi-day risk trajectories for all grid cells.

SECTION 9

Risk Scoring Methodology

INFERNIS assigns each grid cell a composite risk score mapped to a six-level danger classification system: VERY_LOW, LOW, MODERATE, HIGH, VERY_HIGH, and EXTREME. This classification aligns with the familiar CFFDRS danger scale while incorporating ML-derived probability refinement.

Raw model outputs undergo probability calibration via Platt scaling (logistic regression on a held-out calibration set) to ensure that a predicted probability of 0.30 corresponds to an actual observed fire frequency of approximately 30% in similar conditions. This calibration is performed independently for each BEC zone to account for the dramatically different baseline fire rates across the province – the Coastal Western Hemlock zone has a fundamentally different fire regime than the Interior Douglas-fir or Boreal White and Black Spruce zones.

The composite score integrates calibrated model outputs weighted by zone-specific coefficients determined through logistic regression on historical data. In the current 1km calibration, XGBoost probability dominates the composite. Danger class thresholds are set to optimize the tradeoff between false alarm rate and detection probability, with higher sensitivity in zones protecting communities and critical infrastructure.

LEVEL	SCORE RANGE	DESCRIPTION
VERY LOW	0.00 – 0.05	Negligible risk. Wet or snow-covered conditions.
LOW	0.05 – 0.15	Minor risk. Fires unlikely under current conditions.
MODERATE	0.15 – 0.35	Elevated. Fires possible with an ignition source.
HIGH	0.35 – 0.60	Significant. Fires likely to spread if ignited.
VERY HIGH	0.60 – 0.80	Severe. Aggressive fire behavior expected.
EXTREME	0.80 – 1.00	Critical. Explosive fire growth potential.

SECTION 10

Use Cases

Government and Fire Services

Pre-positioning suppression crews and equipment based on next-day risk predictions. Proactive evacuation planning for communities in forecast high-risk zones. Optimized allocation of limited aerial suppression resources across multiple simultaneous fire starts.

Insurance

Wildfire risk underwriting at the property level, informed by location-specific historical and predicted fire probability. Portfolio exposure analysis across insured properties in BC. Dynamic risk assessment for parametric wildfire insurance products.

Utilities

BC Hydro transmission corridor risk assessment for preventive de-energization decisions. Vegetation management prioritization along power line rights-of-way. Infrastructure hardening investment planning based on long-term risk trends.

Forestry

Harvest scheduling that accounts for predicted fire risk windows. Fire guard construction and maintenance prioritization. Reforestation site selection informed by projected fire frequency.

Public and Developers

Community-level fire risk dashboards for municipal governments. Fire risk data layers for weather applications and outdoor recreation platforms. Integration into mapping tools and real-time alerting systems.

SECTION 11

Competitive Landscape

SYSTEM	COVERAGE	RESOLUTION	ML-ENHANCED	FORECAST	API ACCESS	BC-OPTIMIZED
CWFIS	Canada	Station-based	No	Yes (FWI/FBP grids)	OGC GIS services (no REST)	No
Technosylva Wildfire Analyst	Primarily US	High	Partial	Yes	Enterprise only	No
Ambee Wildfire API	Global	Coarse	Limited	Limited	Yes	No
INFERNIS	**BC**	**1km**	**Yes**	**10-day**	**Yes (REST/JSON)**	**Yes**

The Canadian Wildland Fire Information System (CWFIS) is the incumbent government system. It is scientifically authoritative, publishes gridded FWI maps and multi-day fire danger products, and exposes data layers via OGC web services (WMS/WFS/WCS). However, it is not ML-enhanced and does not provide a modern JSON-based developer API accessible to typical web and application developers. Technosylva's Wildfire Analyst is a sophisticated enterprise platform used by US fire agencies, but it is US-focused, carries significant licensing costs, and is not optimized for Canadian data sources or BC-specific conditions. Ambee provides a global wildfire API but operates at coarse resolution without regional calibration, making it poorly suited for BC-specific risk assessment.

INFERNIS occupies a distinct position: BC-specific optimization with per-BEC-zone calibration, ML-enhanced prediction built on the FWI scientific foundation, 1km spatial resolution, exclusive use of open data sources, and a REST/JSON API designed for diverse consumers from government agencies to mobile application developers.

SECTION 12

Roadmap

Phase 1 – MVP (Complete)

XGBoost occurrence prediction on a 5km grid covering 84,535 cells across all of British Columbia. 24-feature model trained on 10 years of data. Core REST API with geographic query support. Firebase-authenticated self-service dashboard with API key provisioning, usage tracking, and tiered access. Automated daily prediction pipeline from data ingestion through model inference to API delivery.

Phase 2 – Spatial Intelligence (Complete)

FireUNet CNN (7.8M parameters) generating continuous spatial risk heatmaps from 12-channel raster inputs. Per-BEC-zone calibration via logistic regression in logit space across all 14 biogeoclimatic zones. Risk Fuser ensemble combining XGBoost point predictions with CNN spatial context in a weighted logit-space architecture. Feature set expanded from 24 to 28 features: four-depth soil moisture profiles (ERA5-Land), leaf area index (LAI), distance-to-nearest-road (BC Digital Road Atlas), and enhanced vegetation indices.

Phase 3 – High Resolution & Forecasting (Complete)

1km grid generator producing 2,113,524 cells in BC Albers (EPSG:3005) projection with vectorized BEC zone assignment via spatial join. Multi-day forecast pipeline combining HRDPS (days 1–2) and GDPS (days 3–10) NWP data with FWI roll-forward for up to 10-day fire risk trajectories, with confidence decay and per-cell forecast storage. Vectorized data processing pipeline: raster sampling via numpy fancy indexing and ERA5 interpolation via scipy RegularGridInterpolator, enabling efficient feature computation at 2M-cell scale. Chunked parquet output (weekly chunks) with float16 storage for managing multi-gigabyte feature matrices. Walk-forward historical backtesting framework with temporal cross-validation (AUC 0.90–0.93 across 2019–2024), per-BEC-zone accuracy breakdowns, and model comparison tooling. End-to-end 1km retraining pipeline composing feature processing, XGBoost training (AUC 0.974), BEC calibration (14 zones), and CNN training (AUC 0.815, 24 epochs on Apple MPS) into a single script.

SECTION 13

Technical Specifications

COMPONENT	TECHNOLOGY
Language	Python 3.11+
API Framework	FastAPI
Primary ML Model	XGBoost
Spatial ML Model	PyTorch (U-Net)
Database	PostgreSQL with PostGIS
Cache Layer	Redis
Satellite Data Access	Google Earth Engine
Deployment	Railway
Grid Resolution	1km (primary), 5km (legacy)
Grid Coverage	2,113,524 cells at 1km; 84,107 cells at 5km
Forecast Horizon	Up to 10 days (HRDPS days 1–2, GDPS days 3–10)
Backtesting	Walk-forward temporal CV (AUC 0.90–0.93), per-BEC-zone breakdowns
Training Corpus	298,606 labeled samples (10:1 neg ratio), 2015–2025
Model Features	28 (6 FWI + 10 weather + 3 vegetation + 5 topo/infrastructure + 2 temporal + 2 lightning)
CNN Architecture	FireUNet, 12 input channels, 7.8M parameters, 256×512 raster, AUC 0.815
Regional Calibration	Per-BEC-zone logistic regression across 14 zones
XGBoost AUC-ROC	0.974
Average Precision	0.794
Brier Score	0.036
Authentication	Firebase Auth (Google + email/password)
Prediction Frequency	Daily at 14:00 PT + 10-day forecast
API Format	REST/JSON

SECTION 14

Limitations and Uncertainties

INFERNIS is a research-grade system transitioning toward operational deployment. The following limitations should inform interpretation of its outputs and assessment of its current capabilities.

Spatial Resolution vs. Input Resolution

INFERNIS produces predictions at 1km grid resolution, but several key input variables originate at coarser native resolutions. ERA5 reanalysis weather data is natively ~31km; ERA5-Land soil moisture is ~9km. These are bilinearly interpolated to the 1km prediction grid, meaning neighboring cells share substantially similar weather and moisture inputs. The 1km resolution is genuine for topographic features (CDEM at ~23m), road distance, and the prediction grid itself, but weather-driven features do not carry independent information at the 1km scale. Users should interpret the 1km output as topographically and vegetatively refined risk estimates built on ~10–30km weather inputs, not as independently measured conditions at each kilometer.

Ensemble Weighting

The hybrid architecture includes both XGBoost and CNN (FireUNet) models, combined via per-zone logistic regression calibration. In current calibration on the 1km grid, the CNN receives near-zero weight across all 14 BEC zones – the ensemble is effectively XGBoost-only. This indicates that at 1km resolution with the current CNN architecture and training data, the cell-level XGBoost model already captures the predictive signal that the CNN's spatial context would provide. The CNN architecture remains in the pipeline and may provide additional value with higher-resolution spatial inputs, more expressive training, or alternative fusion strategies. The reported AUC-ROC of 0.974 reflects the XGBoost component; the CNN's independent AUC of 0.815 contributes minimally to the combined output.

Feature Utilization

The feature vector is designed for 28 inputs, but SHAP analysis reveals that several features contribute zero marginal predictive value in the current 1km model: 24-hour precipitation, evapotranspiration, slope, aspect, hillshade, and both lightning density features all show zero mean |SHAP| values. Their information may be captured by correlated features (e.g., precipitation effects absorbed by FWI moisture codes; slope/aspect absorbed by elevation). The zero lightning SHAP is notable given that lightning is responsible for roughly 60% of BC wildfire ignitions – this likely reflects limitations in the available lightning data resolution or temporal alignment rather than the irrelevance of lightning as a fire driver. XGBoost's native gain-based importance does assign non-zero values to these features, indicating they contribute to some splits but not materially to overall prediction quality.

Label Definition

Positive labels are derived from historical fire records (CNFDB point-of-origin and BC Wildfire perimeters) with a 3km/3-day spatiotemporal buffer. This means the model predicts fire-conducive conditions in a neighborhood, not the precise ignition cell. Small fires (<4 hectares) are underrepresented in the historical record, and reporting consistency varies across regions and decades. The 10:1 negative sampling ratio, while standard in imbalanced classification, means the training distribution differs from the true base rate of fire occurrence (approximately 0.01–0.1% of cells on any given day), and calibrated probabilities should be interpreted accordingly.

Operational Calibration

Cross-validated AUC-ROC (0.974) and walk-forward backtest AUC (0.90–0.93) measure discrimination – the model's ability to rank fire-prone cells above non-fire cells. They do not directly measure calibration accuracy at the extreme low base rates encountered operationally. A Brier score of 0.036 indicates good overall calibration, but Brier scores can be dominated by the large number of true negatives. Reliability at operational decision thresholds (e.g., the top 1% of predictions) should be evaluated with additional metrics such as precision-recall curves and calibration diagrams before deployment in safety-critical applications.

Smoke and Extreme Events

Model performance during active large fire events – when smoke significantly degrades satellite imagery quality (NDVI, LAI) and atmospheric conditions deviate from reanalysis assumptions – has not been independently evaluated. The 2023 season is included in the training data, providing some exposure to extreme conditions, but systematic evaluation of prediction quality under heavy smoke is an area for future work.

Data Licensing

Google Earth Engine is currently used under its non-commercial license for satellite data access. This is appropriate for research and open-source development but must be upgraded to a commercial license before INFERNIS generates revenue from API access. ERA5, CNFDB, CDEM, and other government data sources are freely available for commercial use.

SECTION 15

Conclusion

Wildfire management in British Columbia stands at an inflection point. The old paradigm – reactive response guided by station-interpolated danger ratings – was built for a climate and a landscape that no longer exist. The 2023 season, which burned more hectares than the previous two record-setting years combined, demonstrated that the frequency and intensity of BC wildfires have moved beyond the envelope that legacy systems were designed to handle.

INFERNIS represents the next generation of wildfire prediction: machine learning models trained on the richest open data ecosystem in the world, validated against peer-reviewed research, calibrated to the specific biogeoclimatic diversity of British Columbia, and delivered through a modern API architecture designed for integration into the systems that governments, insurers, utilities, and communities depend on.

Canada publishes more open environmental data than nearly any nation on Earth. ERA5 reanalysis provides gap-free gridded weather back to 1940. The CNFDB catalogs decades of fire history. MODIS and Sentinel-2 observe the province daily from orbit. This data exists. The science to turn it into predictions has been published and validated. What has been missing is an engineered system that assembles these inputs, applies modern ML, and delivers actionable intelligence to the people and organizations who need it.

INFERNIS is that system – built on open data, validated through rigorous backtesting, and delivered as a modern API. Intelligence forged in fire.

References

Intelligence forged in fire.

[1] BCWildfire: A Long-term Multi-factor Dataset and Deep Learning Benchmark for Boreal Wildfire Risk Prediction. arXiv:2511.17597, 2025.

[2] Province of British Columbia, "What causes wildfire." https://www2.gov.bc.ca/gov/content/safety/wildfire-status/wildfire-response/what-causes-wildfire

[3] Natural Resources Canada / Environment and Climate Change Canada, "Lightning and forest fires." https://www.canada.ca/en/environment-climate-change/services/lightning/forest-fires.html

[4] 2017 British Columbia wildfires. https://en.wikipedia.org/wiki/2017_British_Columbia_wildfires

[5] 2018 British Columbia wildfires. https://en.wikipedia.org/wiki/2018_British_Columbia_wildfires

[6] 2023 Canadian wildfires. https://en.wikipedia.org/wiki/2023_Canadian_wildfires

[7] Statistics confirm devastating 2023 wildfire season. https://www.ubcm.ca/about-ubcm/latest-news/statistics-confirm-devastating-2023-wildfire-season

[8] CIFFC 2023 Canada Report. https://ciffc.ca/sites/default/files/2024-03/03.07.24_CIFFC_2023CanadaReport%20(1).pdf

[9] Canadian Forest Fire Weather Index (FWI) System. https://cwfis.cfs.nrcan.gc.ca/background/summary/fwi

[10] Canada’s Fire Weather Index System. https://natural-resources.canada.ca/forest-forestry/wildland-fires/canada-fire-weather-index-system

[11] CWFIS Data Services. https://cwfis.cfs.nrcan.gc.ca/downloads/CWFIS_DataServices_HowtoAccessDailyMaps&DataLayers.pdf

[12] CWFIS Datamart. https://cwfis.cfs.nrcan.gc.ca/datamart

[13] ERA5 atmospheric reanalysis. https://cds.climate.copernicus.eu/datasets/reanalysis-era5-complete

[14] ERA5-Land hourly data from 1950 to present. https://cds.climate.copernicus.eu/datasets/reanalysis-era5-land

[15] Canadian National Fire Database (CNFDB). https://cwfis.cfs.nrcan.gc.ca/ha/nfdb

[16] Canadian Lightning Detection Network. https://www.canada.ca/en/environment-climate-change/services/lightning/canadian-detection-network.html

INFERNIS

INTELLIGENCE FORGED IN FIRE

A Machine Learning Engine for Wildfire Prediction in British Columbia

White Paper — Version 1.0

Date	February 2026
Author	Arnold Dogelis
Contact	hello@argonbi.com
Location	Vancouver, BC

Table of Contents

1. Executive Summary
2. The Problem: Wildfire Risk in British Columbia
3. The Solution: INFERNIS
4. The Science
5. Empirical Feature Importance
6. Research Validation
7. Data Foundation
8. System Architecture
9. Risk Scoring Methodology
10. Use Cases
11. Competitive Landscape
12. Roadmap
13. Technical Specifications
14. Limitations and Uncertainties
15. Conclusion
16. References

SECTION 1

Executive Summary

INFERNIS is a machine learning engine purpose-built to predict wildfire occurrence and spatial risk across British Columbia. By combining gradient-boosted decision trees with convolutional neural networks in a regionally calibrated ensemble, INFERNIS ingests 21 open data sources – spanning satellite imagery, reanalysis weather grids, multi-depth soil moisture profiles, vegetation indices, and historical fire records – to produce calibrated daily fire occurrence probabilities at 1km grid resolution across the entire province.

The problem INFERNIS addresses is both urgent and structurally underserved. British Columbia has experienced three of its worst wildfire seasons in recorded history within the last decade, with cumulative damages exceeding tens of billions of dollars and hundreds of thousands of residents displaced. Existing fire danger systems, while scientifically rigorous, remain station-based and reactive. They were not designed to deliver the granular, probabilistic, machine-learning-enhanced predictions that modern fire management, insurance underwriting, and infrastructure planning demand. INFERNIS closes that gap: it transforms Canada’s world-class open data ecosystem into actionable, API-delivered wildfire intelligence.

Key differentiators include BC-specific model calibration across 14 biogeoclimatic zones, a hybrid ensemble architecture validated through walk-forward temporal backtesting, automated daily ingestion of open data sources requiring zero manual intervention, and a REST API designed for integration by government agencies, insurers, utilities, and application developers. INFERNIS is built entirely on freely available Canadian and international open data.

SECTION 2

The Problem: Wildfire Risk in British Columbia

British Columbia is in a wildfire crisis that is accelerating, not stabilizing. The 2017 fire season burned 1.2 million hectares and was declared the worst in provincial history. That record fell the very next year when 2018 saw 1.35 million hectares consumed. Both were eclipsed by the 2023 season, which burned 2.84 million hectares – more than double the previous record and the worst wildfire season in Canadian history. This is not cyclical variation. It is a structural shift driven by climate change, fuel accumulation from decades of fire suppression, and expanding human-wildland interface.

The economic toll is staggering. The 2023 season is expected to have caused multi-billion-dollar losses when accounting for direct suppression expenditures (over $1 billion CAD in BC alone), insured property losses, infrastructure damage, public health impacts from smoke exposure, and economic disruption from evacuations that displaced over 35,000 residents. Insurance losses from wildfire are now a material line item in Canadian reinsurance portfolios, and multiple insurers have begun restricting coverage in high-risk BC communities.

Current wildfire danger assessment in Canada relies on the Canadian Forest Fire Danger Rating System (CFFDRS), a scientifically rigorous framework developed over decades by the Canadian Forest Service. The CFFDRS and its Fire Weather Index (FWI) System remain the gold standard for fire danger rating worldwide. However, the system was designed for an era of manual weather station readings and human interpretation. It is historically station-based, though gridded FWI products based on ERA5-style inputs now exist. The core operational system does not incorporate machine learning, satellite-derived vegetation indices, or multi-depth soil moisture reanalysis. The Canadian Wildland Fire Information System (CWFIS) does publish daily and forecast fire danger maps and exposes data layers via OGC web services (WMS/WFS/WCS), but it does not provide a modern JSON-based developer API. Its documentation and interfaces are oriented toward government fire management workflows and GIS professionals rather than commercial developers, insurers, or application builders.

The gap between what the science makes possible and what the operational systems deliver represents both a public safety risk and a commercial opportunity.

SECTION 3

The Solution: INFERNIS

INFERNIS is a machine learning-powered fire prediction engine designed specifically for British Columbia. It operates as a daily batch prediction system, ingesting data each afternoon after noon weather observations are finalized, and producing grid-level fire occurrence probabilities and danger classifications that are served via a REST API.

The architecture employs a hybrid ensemble approach. An XGBoost gradient-boosted classifier serves as the primary occurrence prediction model, trained on approximately 298,000 labeled samples (27,146 positive, 271,460 negative at a 10:1 ratio with spatiotemporal buffering) spanning eleven fire seasons (2015–2025) and incorporating 28 engineered features drawn from weather reanalysis, FWI components, multi-depth soil moisture, vegetation indices, topographic derivatives, infrastructure proximity, and lightning detection. In cross-validated evaluation on 1km grid data, the XGBoost model achieves an AUC-ROC of 0.974 – meaning it correctly ranks fire-prone conditions above non-fire conditions 97.4% of the time. A FireUNet convolutional neural network (7.8M parameters) operates in parallel, processing 12-channel spatial inputs across a 256×512 raster grid to generate continuous risk heatmaps that capture landscape-scale fire spread patterns and spatial autocorrelation that point-based models miss, achieving an AUC-ROC of 0.815. The Risk Fuser combines both model outputs in logit space with per-zone calibration coefficients independently tuned for each of BC's 14 biogeoclimatic (BEC) zones, producing a final composite risk score for each grid cell.

Beyond current-day predictions, INFERNIS produces multi-day fire risk forecasts by ingesting numerical weather prediction (NWP) data from Environment and Climate Change Canada. High-Resolution Deterministic Prediction System (HRDPS) forecasts at 2.5km resolution drive the first two forecast days, while Global Deterministic Prediction System (GDPS) data extends predictions out to 10 days. FWI moisture codes are rolled forward day-by-day using forecast weather, maintaining physical consistency across the forecast horizon. A confidence decay factor attenuates predictions at longer lead times to transparently communicate increasing uncertainty.

All data sources are open and freely available from Canadian federal agencies, provincial data portals, and international scientific repositories. INFERNIS requires no proprietary data subscriptions. Google Earth Engine is used under its non-commercial license for satellite data access during development, with a plan to transition to a commercial license once revenue is generated. The system is designed for full automation: once configured, the daily pipeline runs without human intervention, from data retrieval through model inference to API delivery.

SECTION 4

The Science

The Fire Weather Index System

INFERNIS builds on the scientific foundation of the Canadian Fire Weather Index (FWI) System, which models fire danger through a three-tier structure of moisture codes and fire behavior indices.

At the base tier, three moisture codes track fuel drying at different time scales. The Fine Fuel Moisture Code (FFMC) represents the moisture content of surface litter and fine fuels, responding to weather changes within hours. The Duff Moisture Code (DMC) tracks moisture in loosely compacted organic layers, operating on a time scale of days to weeks. The Drought Code (DC) models deep organic soil moisture with a seasonal memory spanning weeks to months. Together, these three codes encode cumulative drying across the full spectrum of fuel layers relevant to fire ignition and behavior.

The intermediate tier combines these moisture codes into two compound indices. The Initial Spread Index (ISI) merges FFMC with wind speed to estimate the expected rate of fire spread. The Buildup Index (BUI) combines DMC and DC to represent the total fuel available for combustion.

At the top tier, the Fire Weather Index (FWI) integrates ISI and BUI into a single numeric rating of fire intensity. INFERNIS uses all six FWI components as engineered features in its ML models, preserving the decades of fire science encoded in their formulations while allowing the machine learning layer to discover nonlinear interactions and threshold effects that the linear FWI aggregation cannot capture.

Machine Learning Enhancement

Beyond the classical FWI features, INFERNIS incorporates modern data sources that were unavailable when the CFFDRS was designed. Satellite-derived vegetation indices (NDVI, EVI, LAI) from MODIS and Sentinel-2 provide direct observation of fuel condition, canopy structure, and vegetation stress. ERA5 reanalysis provides gridded, gap-free precipitation and evapotranspiration measurements alongside soil moisture at four depths spanning from the surface through to deep soil layers. High-resolution topographic features derived from the Canadian Digital Elevation Model (CDEM) – elevation, slope, aspect, and hillshade – encode terrain characteristics that influence fire behavior at the landscape scale. Distance to the nearest road, derived from the BC Digital Road Atlas, captures human-wildland interface proximity, a key factor in both ignition likelihood and suppression access. Lightning detection from the Canadian Lightning Detection Network (CLDN) captures the primary natural ignition source – lightning is responsible for roughly 60% of BC wildfire ignitions and, as in the rest of Canada, for the majority of the total area burned.

The complete 28-feature vector per grid cell per day comprises: 6 FWI components (FFMC, DMC, DC, ISI, BUI, FWI), 10 weather variables (temperature, relative humidity, wind speed, wind direction, 24h precipitation, soil moisture at four depths, evapotranspiration), 3 vegetation indices (NDVI, snow cover fraction, leaf area index), 5 topographic and infrastructure features (elevation, slope, aspect, hillshade, distance to nearest road), 2 temporal encodings (day-of-year sine and cosine), and 2 lightning features (24h and 72h flash density).

Training Data Construction

Labels are derived from the Canadian National Fire Database (CNFDB) point-of-origin records and BC Wildfire Service perimeter data. A grid cell is labeled positive for a given day if a fire ignition point falls within its 1km boundary during that day, or if the cell is within a 3km spatial buffer and 3-day temporal window of an ignition point. This buffering reflects the reality that fires do not ignite instantaneously at a single point – ignition conditions exist across a spatial neighborhood before and during the event. Negative samples are drawn from fire-free cells with spatiotemporal exclusion: negatives must be at least 10km and 7 days from any fire event, preventing contamination from near-miss conditions that are functionally fire-prone. The negative-to-positive ratio is 10:1 (271,460 negatives to 27,146 positives in the 1km training set), sampled with stratification across years and BEC zones to prevent temporal or geographic bias from dominating the training signal.

This label construction means that INFERNIS predicts fire occurrence conditions (whether a cell exhibits the combination of weather, fuel, and landscape factors associated with nearby ignition) rather than predicting the precise location of a specific fire start to the exact grid cell. The spatiotemporal buffering also serves as a data leakage prevention mechanism: by excluding near-fire cells from the negative pool, the model cannot learn trivially from spatial autocorrelation of adjacent fire/non-fire cells.

SECTION 5

Empirical Feature Importance

Training on the 1km grid across eleven BC fire seasons with the full 28-feature vector reveals the following feature importance ranking by mean |SHAP| value:

RANK	FEATURE	MEAN SHAP	CATEGORY
1	NDVI (vegetation greenness)	1.25	Vegetation
2	Elevation	1.03	Topography
3	DMC (duff moisture code)	0.88	FWI
4	DC (drought code)	0.74	FWI
5	Soil moisture	0.49	Weather
6	Day-of-year	0.44	Temporal
7	FFMC (fine fuel moisture code)	0.42	FWI
8	Temperature	0.31	Weather
9	Wind speed	0.28	Weather

Vegetation condition (NDVI) remains the single most important predictor. Elevation ranks second, demonstrating that topographic context materially improves prediction at 1km resolution. Three FWI components (DMC, DC, FFMC) appear in the top nine, confirming the value of the classical fire weather indices as ML features. Soil moisture ranks fifth, capturing landscape-level dryness that integrates weeks of precipitation and evapotranspiration history.

SECTION 6

Research Validation

The approach is validated by a growing body of peer-reviewed research. Recent machine learning studies on regional wildfire prediction using gradient-boosted models with ERA5 and FWI feature sets report AUCs in the 0.8–0.9 range, depending on region, task definition, and evaluation protocol. The BCWildfire benchmark dataset, which evaluates deep learning models at 1km resolution across British Columbia, reports that recent architectures (e.g., S-Mamba) achieve F1 scores above 0.85 and PR-AUC close to 0.95 in boreal wildfire risk prediction.

INFERNIS achieves an AUC-ROC of 0.974 and average precision of 0.794 in 5-fold stratified cross-validation on its 1km, 11-year training corpus, with a Brier score of 0.036 indicating well-calibrated probability outputs. Walk-forward temporal backtesting (training on years [2015, test_year-1], testing on test_year) yields AUC-ROC values of 0.90–0.93 across six held-out fire seasons (2019–2024), confirming that model performance generalizes across years and is not an artifact of random cross-validation splits.

SECTION 7

Data Foundation

INFERNIS draws from 21 open data sources spanning eight major categories. A comprehensive catalog is maintained in the project's DATA_SOURCES.md document; the following summarizes the key inputs.

Historical Fires

The Canadian National Fire Database (CNFDB) provides point-of-origin records for fires dating back decades. BC Wildfire Service perimeter data supplies polygon boundaries for all significant fires, enabling both point-based classification training and spatial burn-area modeling.

Weather

ERA5 reanalysis from the European Centre for Medium-Range Weather Forecasts (ECMWF) serves as the primary weather backbone. ERA5 provides hourly, gridded, gap-free atmospheric variables at approximately 31km native resolution, globally, from 1940 to present with a five-day lag. INFERNIS ingests 2m temperature, dewpoint, 10m wind components, total precipitation, potential evapotranspiration, and soil moisture at multiple depths. Precipitation and evapotranspiration rank among the top 10 most important predictive features, confirming the value of ERA5 variables beyond what the classical FWI system incorporates.

Satellite Imagery

MODIS and Sentinel-2 imagery, accessed via Google Earth Engine, provides vegetation indices (NDVI, EVI), active fire detections (MODIS Thermal Anomalies), and burn severity assessments.

Soil Moisture

ERA5-Land soil moisture layers at four depth levels (0–7cm, 7–28cm, 28–100cm, 100–289cm) provide gridded subsurface water content critical for predicting deep organic fuel drying. All four layers are ingested as model features, capturing moisture gradients from surface litter through to deep root zones.

Vegetation and Fuel

NDVI, EVI, and Leaf Area Index (LAI) time series characterize vegetation phenology, canopy density, and stress. CFFDRS Fuel Behaviour Prediction (FBP) system fuel type maps classify the landscape into standardized fuel categories.

Topography

The Canadian Digital Elevation Model (CDEM) at approximately 23m resolution provides elevation, from which slope gradient, aspect angle, and hillshade illumination index are derived via numerical gradient computation. Elevation ranks as the 2nd most important feature in the trained 1km model (mean |SHAP| = 1.03), demonstrating that topographic context is a critical contributor to fire prediction.

Infrastructure

The BC Digital Road Atlas provides road network geometry used to compute distance-to-nearest-road for each grid cell, capturing human-wildland interface proximity relevant to both ignition probability and suppression accessibility.

Lightning

The Canadian Lightning Detection Network (CLDN) provides lightning strike locations and polarity, as lightning is responsible for roughly 60% of BC wildfire ignitions and, as in the rest of Canada, for the majority of the total area burned.

All data is sourced from Canadian federal agencies, the Government of British Columbia, ECMWF, NASA, and ESA. No proprietary data subscriptions are required.

SECTION 8

System Architecture

INFERNIS is organized into six core subsystems that execute as a coordinated daily pipeline.

DATA FORGE

DATA FORGE is the automated ingestion layer. It retrieves, validates, and standardizes data from all sources on a daily schedule, handling format conversions, coordinate reprojection, temporal alignment, and quality control. Data Forge maintains a local mirror of key datasets and performs incremental updates to minimize bandwidth and processing time.

FIRE CORE

FIRE CORE is the primary prediction engine, built on XGBoost. It operates on a structured 28-feature matrix with one row per grid cell per day, incorporating FWI components, weather variables, multi-depth soil moisture, vegetation indices, topographic features, infrastructure proximity, lightning activity, and temporal encodings. Trained on approximately 298,000 labeled samples (10:1 negative-to-positive ratio with spatiotemporal buffering) spanning 2015–2025 at 1km resolution, the model achieves an AUC-ROC of 0.974 and a Brier score of 0.036 in cross-validated evaluation, producing well-calibrated occurrence probabilities.

HEATMAP ENGINE

HEATMAP ENGINE employs a FireUNet convolutional neural network (7.8M parameters) that processes 12-channel spatial inputs at 256×512 pixel resolution – stacked grids of weather, moisture, vegetation, and topography covering the full BC extent – to generate continuous spatial risk surfaces. The CNN captures landscape-scale patterns, spatial gradients, and neighborhood effects that the cell-independent XGBoost model cannot represent.

RISK FUSER

RISK FUSER combines outputs from FIRE CORE and HEATMAP ENGINE using a weighted ensemble operating in logit space with regional calibration. Calibration coefficients are fitted independently for each of BC's 14 biogeoclimatic (BEC) zones via logistic regression, accounting for dramatically different fire regimes, fuel types, and climate characteristics. The fuser transforms model outputs to logit space, applies zone-specific linear calibration, and maps the resulting probabilities to a six-level danger classification. In current calibration, XGBoost dominates the ensemble weighting across most zones – the CNN spatial risk component receives minimal weight in the logistic regression fit, indicating that the cell-level XGBoost predictions already capture most of the predictive signal at 1km resolution. The CNN architecture remains in the pipeline as an active component; future work on higher-resolution spatial inputs and more expressive CNN training may unlock additional ensemble value.

FORECAST ENGINE

FORECAST ENGINE extends predictions beyond the current day by combining high-resolution HRDPS weather forecasts (days 1–2, 2.5km resolution from Environment and Climate Change Canada) with global GDPS forecasts (days 3–10, 25km resolution). The engine rolls forward FWI moisture codes day-by-day using forecast weather, builds the full 28-feature matrix for each lead day, and applies the XGBoost model to produce multi-day fire risk trajectories. A confidence decay factor (default 0.95 per lead day) attenuates predictions at longer lead times to reflect increasing forecast uncertainty. Forecast weather is bilinearly interpolated from the native NWP grid to the INFERNIS prediction grid using scipy's RegularGridInterpolator.

REST API

REST API serves pre-computed daily predictions and multi-day forecasts via a FastAPI application. Consumers can query fire risk by geographic coordinates, grid cell identifiers, or bounding boxes, receiving JSON responses containing occurrence probabilities, danger classifications, contributing factor breakdowns, and up to 10 days of forecast trajectories with per-day confidence scores and data source attribution (HRDPS or GDPS).

The full pipeline executes daily at 14:00 Pacific Time, after noon weather observations have been incorporated into data feeds. The daily pipeline produces current-day predictions; the forecast pipeline runs immediately afterward, generating multi-day risk trajectories for all grid cells.

SECTION 9

Risk Scoring Methodology

INFERNIS assigns each grid cell a composite risk score mapped to a six-level danger classification system: VERY_LOW, LOW, MODERATE, HIGH, VERY_HIGH, and EXTREME. This classification aligns with the familiar CFFDRS danger scale while incorporating ML-derived probability refinement.

Raw model outputs undergo probability calibration via Platt scaling (logistic regression on a held-out calibration set) to ensure that a predicted probability of 0.30 corresponds to an actual observed fire frequency of approximately 30% in similar conditions. This calibration is performed independently for each BEC zone to account for the dramatically different baseline fire rates across the province – the Coastal Western Hemlock zone has a fundamentally different fire regime than the Interior Douglas-fir or Boreal White and Black Spruce zones.

The composite score integrates calibrated model outputs weighted by zone-specific coefficients determined through logistic regression on historical data. In the current 1km calibration, XGBoost probability dominates the composite. Danger class thresholds are set to optimize the tradeoff between false alarm rate and detection probability, with higher sensitivity in zones protecting communities and critical infrastructure.

LEVEL	SCORE RANGE	DESCRIPTION
VERY LOW	0.00 – 0.05	Negligible risk. Wet or snow-covered conditions.
LOW	0.05 – 0.15	Minor risk. Fires unlikely under current conditions.
MODERATE	0.15 – 0.35	Elevated. Fires possible with an ignition source.
HIGH	0.35 – 0.60	Significant. Fires likely to spread if ignited.
VERY HIGH	0.60 – 0.80	Severe. Aggressive fire behavior expected.
EXTREME	0.80 – 1.00	Critical. Explosive fire growth potential.

SECTION 10

Use Cases

Government and Fire Services

Pre-positioning suppression crews and equipment based on next-day risk predictions. Proactive evacuation planning for communities in forecast high-risk zones. Optimized allocation of limited aerial suppression resources across multiple simultaneous fire starts.

Insurance

Wildfire risk underwriting at the property level, informed by location-specific historical and predicted fire probability. Portfolio exposure analysis across insured properties in BC. Dynamic risk assessment for parametric wildfire insurance products.

Utilities

BC Hydro transmission corridor risk assessment for preventive de-energization decisions. Vegetation management prioritization along power line rights-of-way. Infrastructure hardening investment planning based on long-term risk trends.

Forestry

Harvest scheduling that accounts for predicted fire risk windows. Fire guard construction and maintenance prioritization. Reforestation site selection informed by projected fire frequency.

Public and Developers

Community-level fire risk dashboards for municipal governments. Fire risk data layers for weather applications and outdoor recreation platforms. Integration into mapping tools and real-time alerting systems.

SECTION 11

Competitive Landscape

SYSTEM	COVERAGE	RESOLUTION	ML-ENHANCED	FORECAST	API ACCESS	BC-OPTIMIZED
CWFIS	Canada	Station-based	No	Yes (FWI/FBP grids)	OGC GIS services (no REST)	No
Technosylva Wildfire Analyst	Primarily US	High	Partial	Yes	Enterprise only	No
Ambee Wildfire API	Global	Coarse	Limited	Limited	Yes	No
INFERNIS	**BC**	**1km**	**Yes**	**10-day**	**Yes (REST/JSON)**	**Yes**

The Canadian Wildland Fire Information System (CWFIS) is the incumbent government system. It is scientifically authoritative, publishes gridded FWI maps and multi-day fire danger products, and exposes data layers via OGC web services (WMS/WFS/WCS). However, it is not ML-enhanced and does not provide a modern JSON-based developer API accessible to typical web and application developers. Technosylva's Wildfire Analyst is a sophisticated enterprise platform used by US fire agencies, but it is US-focused, carries significant licensing costs, and is not optimized for Canadian data sources or BC-specific conditions. Ambee provides a global wildfire API but operates at coarse resolution without regional calibration, making it poorly suited for BC-specific risk assessment.

INFERNIS occupies a distinct position: BC-specific optimization with per-BEC-zone calibration, ML-enhanced prediction built on the FWI scientific foundation, 1km spatial resolution, exclusive use of open data sources, and a REST/JSON API designed for diverse consumers from government agencies to mobile application developers.

SECTION 12

Roadmap

Phase 1 – MVP (Complete)

XGBoost occurrence prediction on a 5km grid covering 84,535 cells across all of British Columbia. 24-feature model trained on 10 years of data. Core REST API with geographic query support. Firebase-authenticated self-service dashboard with API key provisioning, usage tracking, and tiered access. Automated daily prediction pipeline from data ingestion through model inference to API delivery.

Phase 2 – Spatial Intelligence (Complete)

FireUNet CNN (7.8M parameters) generating continuous spatial risk heatmaps from 12-channel raster inputs. Per-BEC-zone calibration via logistic regression in logit space across all 14 biogeoclimatic zones. Risk Fuser ensemble combining XGBoost point predictions with CNN spatial context in a weighted logit-space architecture. Feature set expanded from 24 to 28 features: four-depth soil moisture profiles (ERA5-Land), leaf area index (LAI), distance-to-nearest-road (BC Digital Road Atlas), and enhanced vegetation indices.

Phase 3 – High Resolution & Forecasting (Complete)

1km grid generator producing 2,113,524 cells in BC Albers (EPSG:3005) projection with vectorized BEC zone assignment via spatial join. Multi-day forecast pipeline combining HRDPS (days 1–2) and GDPS (days 3–10) NWP data with FWI roll-forward for up to 10-day fire risk trajectories, with confidence decay and per-cell forecast storage. Vectorized data processing pipeline: raster sampling via numpy fancy indexing and ERA5 interpolation via scipy RegularGridInterpolator, enabling efficient feature computation at 2M-cell scale. Chunked parquet output (weekly chunks) with float16 storage for managing multi-gigabyte feature matrices. Walk-forward historical backtesting framework with temporal cross-validation (AUC 0.90–0.93 across 2019–2024), per-BEC-zone accuracy breakdowns, and model comparison tooling. End-to-end 1km retraining pipeline composing feature processing, XGBoost training (AUC 0.974), BEC calibration (14 zones), and CNN training (AUC 0.815, 24 epochs on Apple MPS) into a single script.

SECTION 13

Technical Specifications

COMPONENT	TECHNOLOGY
Language	Python 3.11+
API Framework	FastAPI
Primary ML Model	XGBoost
Spatial ML Model	PyTorch (U-Net)
Database	PostgreSQL with PostGIS
Cache Layer	Redis
Satellite Data Access	Google Earth Engine
Deployment	Railway
Grid Resolution	1km (primary), 5km (legacy)
Grid Coverage	2,113,524 cells at 1km; 84,107 cells at 5km
Forecast Horizon	Up to 10 days (HRDPS days 1–2, GDPS days 3–10)
Backtesting	Walk-forward temporal CV (AUC 0.90–0.93), per-BEC-zone breakdowns
Training Corpus	298,606 labeled samples (10:1 neg ratio), 2015–2025
Model Features	28 (6 FWI + 10 weather + 3 vegetation + 5 topo/infrastructure + 2 temporal + 2 lightning)
CNN Architecture	FireUNet, 12 input channels, 7.8M parameters, 256×512 raster, AUC 0.815
Regional Calibration	Per-BEC-zone logistic regression across 14 zones
XGBoost AUC-ROC	0.974
Average Precision	0.794
Brier Score	0.036
Authentication	Firebase Auth (Google + email/password)
Prediction Frequency	Daily at 14:00 PT + 10-day forecast
API Format	REST/JSON

SECTION 14

Limitations and Uncertainties

INFERNIS is a research-grade system transitioning toward operational deployment. The following limitations should inform interpretation of its outputs and assessment of its current capabilities.

Spatial Resolution vs. Input Resolution

INFERNIS produces predictions at 1km grid resolution, but several key input variables originate at coarser native resolutions. ERA5 reanalysis weather data is natively ~31km; ERA5-Land soil moisture is ~9km. These are bilinearly interpolated to the 1km prediction grid, meaning neighboring cells share substantially similar weather and moisture inputs. The 1km resolution is genuine for topographic features (CDEM at ~23m), road distance, and the prediction grid itself, but weather-driven features do not carry independent information at the 1km scale. Users should interpret the 1km output as topographically and vegetatively refined risk estimates built on ~10–30km weather inputs, not as independently measured conditions at each kilometer.

Ensemble Weighting

The hybrid architecture includes both XGBoost and CNN (FireUNet) models, combined via per-zone logistic regression calibration. In current calibration on the 1km grid, the CNN receives near-zero weight across all 14 BEC zones – the ensemble is effectively XGBoost-only. This indicates that at 1km resolution with the current CNN architecture and training data, the cell-level XGBoost model already captures the predictive signal that the CNN's spatial context would provide. The CNN architecture remains in the pipeline and may provide additional value with higher-resolution spatial inputs, more expressive training, or alternative fusion strategies. The reported AUC-ROC of 0.974 reflects the XGBoost component; the CNN's independent AUC of 0.815 contributes minimally to the combined output.

Feature Utilization

The feature vector is designed for 28 inputs, but SHAP analysis reveals that several features contribute zero marginal predictive value in the current 1km model: 24-hour precipitation, evapotranspiration, slope, aspect, hillshade, and both lightning density features all show zero mean |SHAP| values. Their information may be captured by correlated features (e.g., precipitation effects absorbed by FWI moisture codes; slope/aspect absorbed by elevation). The zero lightning SHAP is notable given that lightning is responsible for roughly 60% of BC wildfire ignitions – this likely reflects limitations in the available lightning data resolution or temporal alignment rather than the irrelevance of lightning as a fire driver. XGBoost's native gain-based importance does assign non-zero values to these features, indicating they contribute to some splits but not materially to overall prediction quality.

Label Definition

Positive labels are derived from historical fire records (CNFDB point-of-origin and BC Wildfire perimeters) with a 3km/3-day spatiotemporal buffer. This means the model predicts fire-conducive conditions in a neighborhood, not the precise ignition cell. Small fires (<4 hectares) are underrepresented in the historical record, and reporting consistency varies across regions and decades. The 10:1 negative sampling ratio, while standard in imbalanced classification, means the training distribution differs from the true base rate of fire occurrence (approximately 0.01–0.1% of cells on any given day), and calibrated probabilities should be interpreted accordingly.

Operational Calibration

Cross-validated AUC-ROC (0.974) and walk-forward backtest AUC (0.90–0.93) measure discrimination – the model's ability to rank fire-prone cells above non-fire cells. They do not directly measure calibration accuracy at the extreme low base rates encountered operationally. A Brier score of 0.036 indicates good overall calibration, but Brier scores can be dominated by the large number of true negatives. Reliability at operational decision thresholds (e.g., the top 1% of predictions) should be evaluated with additional metrics such as precision-recall curves and calibration diagrams before deployment in safety-critical applications.

Smoke and Extreme Events

Model performance during active large fire events – when smoke significantly degrades satellite imagery quality (NDVI, LAI) and atmospheric conditions deviate from reanalysis assumptions – has not been independently evaluated. The 2023 season is included in the training data, providing some exposure to extreme conditions, but systematic evaluation of prediction quality under heavy smoke is an area for future work.

Data Licensing

Google Earth Engine is currently used under its non-commercial license for satellite data access. This is appropriate for research and open-source development but must be upgraded to a commercial license before INFERNIS generates revenue from API access. ERA5, CNFDB, CDEM, and other government data sources are freely available for commercial use.

SECTION 15

Conclusion

Wildfire management in British Columbia stands at an inflection point. The old paradigm – reactive response guided by station-interpolated danger ratings – was built for a climate and a landscape that no longer exist. The 2023 season, which burned more hectares than the previous two record-setting years combined, demonstrated that the frequency and intensity of BC wildfires have moved beyond the envelope that legacy systems were designed to handle.

INFERNIS represents the next generation of wildfire prediction: machine learning models trained on the richest open data ecosystem in the world, validated against peer-reviewed research, calibrated to the specific biogeoclimatic diversity of British Columbia, and delivered through a modern API architecture designed for integration into the systems that governments, insurers, utilities, and communities depend on.

Canada publishes more open environmental data than nearly any nation on Earth. ERA5 reanalysis provides gap-free gridded weather back to 1940. The CNFDB catalogs decades of fire history. MODIS and Sentinel-2 observe the province daily from orbit. This data exists. The science to turn it into predictions has been published and validated. What has been missing is an engineered system that assembles these inputs, applies modern ML, and delivers actionable intelligence to the people and organizations who need it.

INFERNIS is that system – built on open data, validated through rigorous backtesting, and delivered as a modern API. Intelligence forged in fire.

SECTION 16

References

Intelligence forged in fire.

[1] BCWildfire: A Long-term Multi-factor Dataset and Deep Learning Benchmark for Boreal Wildfire Risk Prediction. arXiv:2511.17597, 2025.

[2] Province of British Columbia, "What causes wildfire." https://www2.gov.bc.ca/gov/content/safety/wildfire-status/wildfire-response/what-causes-wildfire

[3] Natural Resources Canada / Environment and Climate Change Canada, "Lightning and forest fires." https://www.canada.ca/en/environment-climate-change/services/lightning/forest-fires.html

[4] 2017 British Columbia wildfires. https://en.wikipedia.org/wiki/2017_British_Columbia_wildfires

[5] 2018 British Columbia wildfires. https://en.wikipedia.org/wiki/2018_British_Columbia_wildfires

[6] 2023 Canadian wildfires. https://en.wikipedia.org/wiki/2023_Canadian_wildfires

[7] Statistics confirm devastating 2023 wildfire season. https://www.ubcm.ca/about-ubcm/latest-news/statistics-confirm-devastating-2023-wildfire-season

[8] CIFFC 2023 Canada Report. https://ciffc.ca/sites/default/files/2024-03/03.07.24_CIFFC_2023CanadaReport%20(1).pdf

[9] Canadian Forest Fire Weather Index (FWI) System. https://cwfis.cfs.nrcan.gc.ca/background/summary/fwi

[10] Canada’s Fire Weather Index System. https://natural-resources.canada.ca/forest-forestry/wildland-fires/canada-fire-weather-index-system

[11] CWFIS Data Services. https://cwfis.cfs.nrcan.gc.ca/downloads/CWFIS_DataServices_HowtoAccessDailyMaps&DataLayers.pdf

[12] CWFIS Datamart. https://cwfis.cfs.nrcan.gc.ca/datamart

[13] ERA5 atmospheric reanalysis. https://cds.climate.copernicus.eu/datasets/reanalysis-era5-complete

[14] ERA5-Land hourly data from 1950 to present. https://cds.climate.copernicus.eu/datasets/reanalysis-era5-land

[15] Canadian National Fire Database (CNFDB). https://cwfis.cfs.nrcan.gc.ca/ha/nfdb

[16] Canadian Lightning Detection Network. https://www.canada.ca/en/environment-climate-change/services/lightning/canadian-detection-network.html

INFERNIS — A Machine Learning Engine for Wildfire Prediction in British Columbia