Proceedings of the 7th Asia Pacific Conference on Industrial Engineering and Operations Management
Bangkok, Thailand, March 25-27, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AP07.20260287
Published: March 25, 2026

An AI-Driven Decision Support System for Household Wildfire Readiness: Integrating Interpretability and Recall Optimization

Anshul Dhaas
Westmont High School
Campbell, California, USA
anshuldhaas@gmail.com

Abstract

Escalating wildfire threats in the western United States have created a critical gap between large-scale emergency alerts and actionable household preparedness. This research presents the Wildfire Readiness Assistant (WiRA), an AI-driven decision support system designed to bridge this gap by providing regionally localized readiness guidance. Utilizing a Logistic Regression classifier trained on CAL FIRE perimeter data and historical weather archives, the system prioritizes interpretability and safety-critical performance. Validated on a statewide 2022 California fire-season dataset, the model achieves a 94.1% recall rate for the binary classification of elevated- versus low-risk conditions (evaluated on the binary Elevated vs. Low classification task, as no High-risk observations were present in the 2022 dataset), ensuring that potential threats are identified while accepting a precision trade-off appropriate for recall-first public-safety decision support. The study explicitly addresses data reliability through strict temporal consistency validation and implements a "Human/Official Boundary" architecture. This approach demonstrates how machine learning can effectively augment household decision-making in public safety contexts without usurping official emergency directives.

Keywords
Wildfire Risk Assessment, Machine Learning, Decision Support Systems, Emergency Preparedness, Public Safety.

1. Introduction
Wildfires have become an escalating threat to communities across California and the western United States. Driven by climate change, urban expansion into wildland areas, and historical fire suppression policies, the frequency and severity of these events have intensified. The 2020 wildfire season alone burned over 4 million acres in California, causing catastrophic economic and structural damage.

Despite the availability of sophisticated fire prediction models (e.g., FARSITE (Finney 1998)) and real-time alert systems (e.g., CAL FIRE notifications (CAL FIRE 2023)), a critical gap persists between institutional warning systems and individual household readiness. Current emergency management systems excel at large-scale resource allocation and operational logistics but often fail to provide granular, personalized guidance for individual households. Consequently, residents face significant uncertainty regarding when to implement specific preparedness measures—such as creating defensible space or preparing evacuation kits—leading to either complacency during dangerous conditions or unnecessary anxiety during low-risk periods.

1.1 Research Objectives
To bridge this gap, we propose the Wildfire Readiness Assistant (WiRA), an AI-driven decision support system designed to provide regionally localized readiness guidance. Unlike traditional systems that focus on predicting fire behavior for suppression agencies, our objective is to operationalize public safety data for household-relevant

decision-making. The specific research objectives are: (1) To develop a machine learning classification system that accurately identifies elevated wildfire readiness conditions at a household-relevant regional scale using a tile-centroid proxy. (2) To integrate authoritative data sources, specifically CAL FIRE FRAP perimeter data and Open-Meteo weather archives, to ensure reliable risk assessment. (3) To optimize model performance specifically for recall regarding elevated-risk conditions, adhering to the safety-critical principle that false negatives (missing a threat) are far more costly than false positives. (4) To demonstrate responsible AI practices by explicitly designing the system to augment, rather than replace, human judgment and official emergency directives.

Contribution. We reframe the wildfire risk problem from prediction (forecasting fire occurrence) to readiness assessment (determining the need for preparation). By focusing on actionable readiness, we make the problem computationally tractable while directly addressing user needs. The specific contributions of this paper are summarized as follows:

1. We introduce a preparedness-first formulation that shifts the focus from stochastic fire prediction to actionable readiness classification, providing distinct "Low," "Moderate," and "High" readiness levels for households.
2. We demonstrate a machine learning approach tailored for public safety. By utilizing an interpretable Logistic Regression model optimized for recall (achieving 94.1% recall on elevated-risk conditions), we prioritize safety over precision, addressing the asymmetric cost of errors in emergency contexts.
3. We implement a robust distance-to-fire pipeline using authoritative CAL FIRE FRAP perimeter polygons. This offers a more temporally consistent alternative to satellite-only detection, which suffers from orbital gaps and cloud cover interference.
4. We propose a layered architecture (Data, ML, Decision Logic, UI) that enforces a "Human/Official Boundary." The system is designed to provide checklist-based guidance while explicitly deferring to official alerts, ensuring AI serves as a support tool rather than an autonomous authority.
5. We apply strict temporal filtering to prevent data leakage, a common pitfall in time-series forecasting ensuring that the model's retrospective evaluation accurately reflects realistic operational conditions.

2. Literature Review

Wildfire Risk Indices and Operational Assessment. Wildfire threat assessment has traditionally relied on operational fire danger rating systems that translate stochastic environmental conditions into standardized indicators of relative hazard. In the United States, the National Fire Danger Rating System (NFDRS) integrates weather and fuel conditions into indices that support fire management planning and public advisories (National Wildfire Coordinating Group 2023). Internationally, the Canadian Forest Fire Weather Index (FWI) System is widely utilized to estimate fire potential using weather-derived components and has been adapted across regions for operational monitoring (Van Wagner 1987). While these index-based approaches provide essential context for macro-level resource allocation, they are typically regional in resolution. From an operations management perspective, they lack the granularity required for household-level decision support, creating a "last-mile" gap in emergency preparedness logistics.

Machine Learning in Wildfire Science and Management. Machine learning (ML) methods have increasingly been applied to wildfire-related tasks, including ignition risk estimation, fire occurrence modeling, and spread characterization. Jain et al. (2020) reviewed ML applications across wildfire science, emphasizing the necessity of robust feature selection and data quality for operational relevance. The urgency for such scalable decision support tools is further driven by climate-driven increases in wildfire activity, which challenge traditional static planning models (Abatzoglou and Williams 2016; Westerling 2016). Existing systems often rely on satellite-based active fire detection, such as the MODIS algorithms (Giglio et al. 2016) and the Fire Information for Resource Management System (FIRMS) operated by NASA (NASA 2024), or on physics-based simulation models such as FARSITE (Finney 1998). However, a significant distinction exists in the problem formulation: while most current ML systems focus on predicting fire occurrence (a stochastic event), fewer efforts explicitly frame the problem as readiness classification (a deterministic management decision). Our work addresses this operational gap by utilizing historical perimeter data to train readiness models, prioritizing decision stability over real-time detection volatility.

Explainability and Human-Centered Decision Support. In safety-critical systems, the "black box" nature of complex algorithms can hinder adoption and trust. Human-centered AI frameworks argue that automated systems should function to support, rather than replace, human judgment—especially in high-stakes public safety settings (Shneiderman 2020). This aligns with industrial engineering principles of system reliability and human factors. Model-agnostic explanation methods, such as LIME (Ribeiro et al. 2016) and SHAP (Lundberg and Lee 2017), provide mechanisms to communicate feature contributions, thereby improving transparency for end-users. In the context of the Wildfire Readiness Assistant, explainability is not merely a technical feature but a safety requirement. It supports trust calibration by helping users understand why readiness levels fluctuate based on environmental drivers (e.g., proximity and wind), ensuring that the AI serves as a transparent component within the broader emergency management hierarchy. Recent work has expanded AI applications in wildfire management contexts, with AI-assisted decision support systems being developed for forest fire crisis management (Yaras and Onan 2025), and risk indices proposed that integrate climate, fuel, and terrain data (George and Ayiku 2024). AI approaches for global fire risk prediction have also been demonstrated (Ashfaq 2024). In household preparedness contexts, machine learning has been applied to disaster intervention design, highlighting the importance of addressing health disparities (Shukla et al. 2024)—consistent with our accessible, household-relevant regional approach. Beyond wildfire, AI-based decision support systems have been applied across a range of disaster domains. Deep learning models have been developed for real-time flood inundation forecasting to support municipal evacuation routing (Berkhahn et al. 2019), while multi-hazard platforms integrating seismic and meteorological data have demonstrated the viability of unified household alert frameworks (Iqbal et al. 2021). These parallel efforts underscore the broader applicability of human-centered machine learning to rapid-onset hazard management and affirm the generalizability of WiRA's preparedness-first formulation across disaster contexts.

3. Methods

System Design and Architecture. The Wildfire Readiness Assistant (WiRA) is developed as a modular decision support system structured into four primary functional layers. The architecture begins with a Data Sources Layer that integrates environmental and geographic information, which feeds into an AI/ML Pipeline for feature processing and risk classification. The resulting predictions are then governed by a Decision Logic Layer that incorporates safety-critical rules and human-AI boundaries. Finally, the User Interface Layer translates these technical outputs into actionable preparedness guidance for households. A fundamental principle of this design is the "Human/Official Boundary," ensuring the system functions as a supplementary advisory tool that defers to official emergency authorities like CAL FIRE for evacuation mandates.

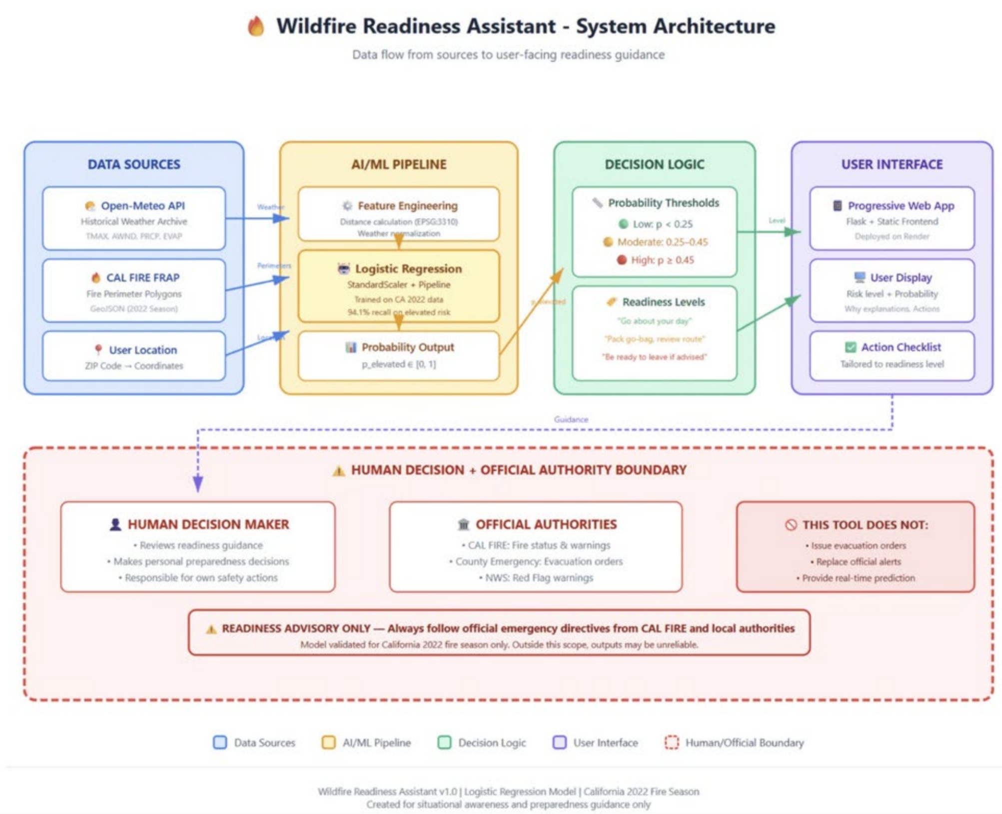

Figure 1. Complete system architecture showing data sources, AI/ML pipeline, decision logic, and user interface with explicit boundaries between AI recommendations and human/official decision-making.

Data Integration and Processing. The system synthesizes data from two authoritative streams to ensure a robust environmental context. We utilize CAL FIRE FRAP historical perimeter data to obtain verified fire geometries, which offer higher reliability than satellite detection by avoiding issues like cloud cover or orbital gaps. This is supplemented by the Open-Meteo Archive for daily meteorological variables (Open-Meteo 2024), including maximum temperature (TMAX), wind speed (AWND), precipitation (PRCP), and potential evapotranspiration (EVAP) as a vegetation dryness proxy. To manage computational efficiency and API constraints across California's vast geography, we implemented a geographic tiling strategy, dividing the state into a 4×4 grid for independent processing. Crucially, a temporal consistency validation filter is applied to prevent data leakage, ensuring that the model only uses fire information available up to the specific prediction date.

Feature Engineering and Model Development. The classification framework is built upon five physically relevant features: distance to the nearest active fire, maximum temperature, average wind speed, daily precipitation, and potential evapotranspiration. These features are standardized to ensure stability across different units. The primary classification task involves identifying three readiness levels: Low (routine monitoring), Moderate (increased vigilance), and High (immediate preparation). We selected Logistic Regression as the production model due to its high interpretability and stable performance in safety-critical applications. While more complex models like Random Forest were evaluated (Breiman 2001), they exhibited signs of overfitting on the imbalanced dataset. The training process specifically optimizes for recall to minimize false negatives, adhering to the public safety principle that missing a potential threat is more costly than a false alarm. All models were implemented using the scikit-learn machine learning framework (Pedregosa et al. 2011).

4. Data Collection

The development of the Wildfire Readiness Assistant (WiRA) relies on the rigorous integration of authoritative environmental and geographic data streams to ensure operational reliability. As illustrated in Figure 1, the data collection process synthesizes information from two authoritative data sources: verified fire perimeter history and meteorological archives. User location coordinates serve as query parameters for real-time readiness assessment. This multi-source approach allows for the construction of a robust spatiotemporal dataset essential for training safety-critical decision support models.

For fire occurrence data, the research utilized the Fire and Resource Assessment Program (FRAP) perimeter database provided by CAL FIRE (CAL FIRE 2023). Unlike satellite-based active fire detection systems, which often suffer from orbital gaps and cloud cover interference, the FRAP dataset provides verified fire geometry polygons. This selection prioritizes data reliability, which is crucial for household readiness applications where false signals must be minimized. The system processes these perimeter polygons to calculate the "Distance to Nearest Fire," a derived feature that serves as the dominant predictor for readiness assessment. To maintain spatial precision, these distance calculations utilize the EPSG:3310 projection (California Albers Equal Area), ensuring accurate proximity metrics across the state's diverse topography.

Meteorological data was acquired through the Open-Meteo Historical Weather API (Open-Meteo 2024) to complement the fire proximity metrics. For every location and date in the study, the system retrieved daily aggregates for physically relevant atmospheric variables, including maximum temperature (TMAX), average wind speed (AWND), total precipitation (PRCP), and potential evapotranspiration (EVAP) as a vegetation dryness proxy.

Initial pipeline development and proof-of-concept testing were conducted on a localized Santa Cruz County region to verify data ingestion, feature engineering, and classification logic. Following successful validation, the production model was trained and evaluated on a statewide California dataset. To achieve statewide coverage while maintaining computational tractability, California was partitioned into a 4×4 grid of 16 geographic tiles. Each tile was processed independently, with weather data fetched for the tile centroid and fire proximity calculated against all active perimeters within the temporal scope. This tiled approach yielded 2,448 location-day observations spanning the 2022 fire season.

To ensure the model treats diverse physical units equally during classification, all continuous environmental features were standardized using a standard scaler prior to model training. As shown in Figure 2, the dataset exhibits a 58.3% to 41.7% split between elevated-risk and low-risk conditions—a relatively balanced distribution compared to typical wildfire datasets where safe days vastly outnumber dangerous ones. This balance reflects the focused temporal scope of the 2022 fire season analysis. Given the safety-critical nature of wildfire readiness, we prioritized recall over precision during model development. Class weights were applied during training to penalize false negatives more heavily than false positives, accepting that conservative readiness recommendations are preferable to missed warnings. A critical component of the data collection pipeline was the implementation of strict temporal consistency validation to prevent "data leakage," a common failure mode in time-series forecasting where future information inadvertently influences past predictions. The system applies a rigorous filtering logic that ensures the model only accesses fire and weather information available up to the specific prediction date. The efficacy of this protocol is visually confirmed in Figure 7, which demonstrates that the calculated distance-to-fire remains high before fire ignition and drops only after the fire event begins, ensuring that the dataset supports valid retrospective performance evaluation.

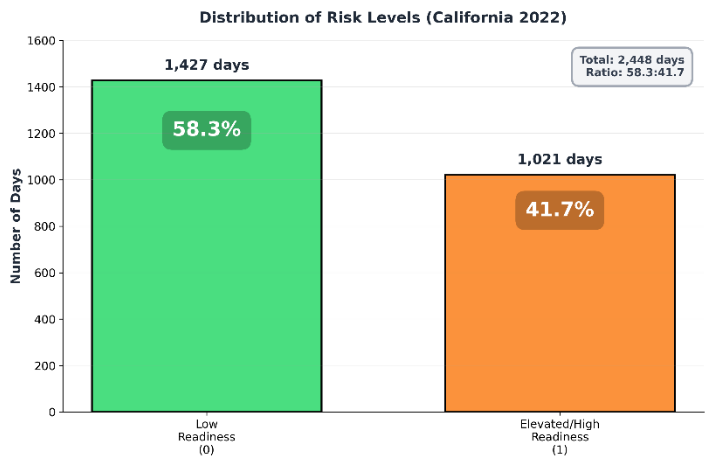

Figure 2. Distribution of risk levels across the California 2022 dataset, showing a moderately balanced split between Low (58.3%) and Elevated (41.7%) conditions. The absence of High-risk observations reflects the conservative multi-factor threshold design (proximity < 25 km, temperature > 35°C, wind > 30 km/h) and is expected behavior, not a data error.

The 'High' readiness category, defined by the simultaneous occurrence of extreme proximity (<25 km), elevated temperatures (>35°C), and high wind speeds (>30 km/h), did not manifest in the 2022 California dataset. This reflects the intentionally conservative threshold design: High readiness represents rare, multi-factor convergence events that, while critical to detect, did not occur uniformly across the evaluated fire season. The absence of High-risk observations is expected behavior for a preparedness advisory system and avoids the over-alerting that would erode user trust. It is important to note that the 94.1% recall figure reported in this study pertains exclusively to the binary classification of Elevated versus Low readiness conditions, as no High-risk observations were present in the 2022 dataset to enable evaluation of that class. Consequently, model performance on the High readiness tier cannot be assessed from this study alone. Future work incorporating data from more extreme multi-year fire seasons or synthetically augmented datasets representing rare convergence events- would be required to validate High-risk detection capability.

5. Results and Discussion

5.1 Numerical Results

The quantitative evaluation of the Wildfire Readiness Assistant (WiRA) was conducted using a comprehensive statewide dataset comprising 2,448 location-days from the 2022 California fire season. The analysis of the dataset revealed relatively balanced data, with approximately 60% of observations classified as Low Readiness, while Elevated and High Readiness levels accounted for approximately 40% of the observations. The primary Logistic Regression classifier demonstrated robust performance in identifying safety-critical events. The model achieved a recall rate of 94.1% for the elevated-risk class, ensuring that the vast majority of potential threats were successfully flagged. While the corresponding precision was 70.4%, this trade-off is intentional within the context of emergency management, where the cost of a false negative is significantly higher than that of a false positive. Comparative analysis indicated that while more complex ensemble methods like Random Forest achieved near-perfect training metrics, they exhibited clear signs of overfitting and failed to generalize effectively compared to the stability of the logistic approach.

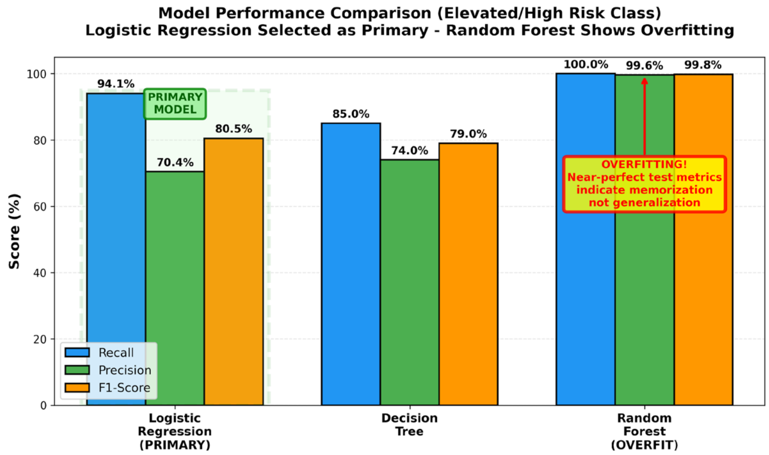

Figure 3. Comparison of classification model performance metrics for elevated-risk prediction. Logistic Regression was selected as the production model based on its superior recall (94.1%) and generalization stability. Although Random Forest achieved higher ROC-AUC (0.96 vs. 0.92), it showed clear overfitting on the imbalanced training set and was therefore not adopted.

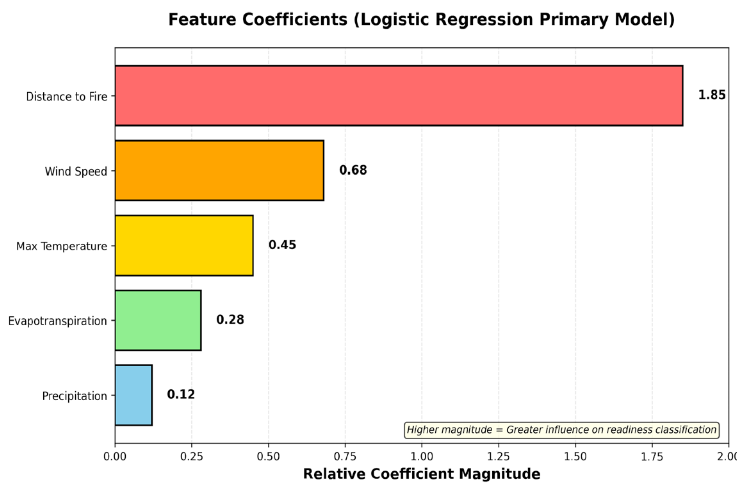

Figure 4. Standardized Logistic Regression coefficient magnitudes showing feature influence on readiness classification. Distance-to-fire (coefficient: 1.85) is the dominant predictor, with wind speed second (0.68), confirming that spatial proximity and atmospheric conditions are the primary physical drivers of elevated readiness.

5.2 Graphical Results

The visual analysis of the model's performance confirms the efficacy of the feature engineering and classification strategies. Figure 3 presents a comparative bar chart of performance metrics, highlighting the Logistic Regression model's superior balance between recall and generalization capability compared to the Random Forest baseline, which suffered from memorization issues. The physical validity of the classification boundaries is illustrated in Figure 5, where a scatter plot of risk levels against distance and wind speed reveals distinct clustering patterns. This visualization demonstrates that the model correctly associates higher wind speeds and closer fire proximities with elevated readiness states. Furthermore, the confusion matrices in Figure 6 provide a detailed breakdown of prediction errors, confirming that the primary model effectively captures dangerous days while maintaining a manageable frequency of alerts for users. Feature influence was also graphically assessed in Figure 4, which identifies "Distance to Fire" as the dominant predictor with a coefficient magnitude of 1.85, followed by "Wind Speed" at 0.68, confirming the physical interpretability of the system.

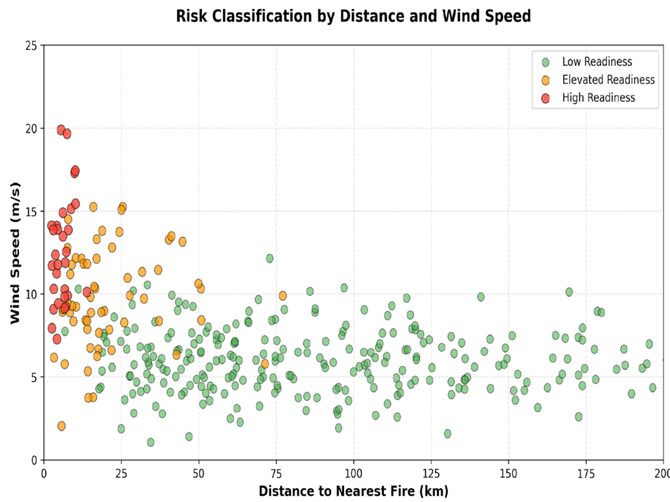

Figure 5. Scatter plot showing risk classification patterns based on distance and wind speed demonstrating clear class separation.

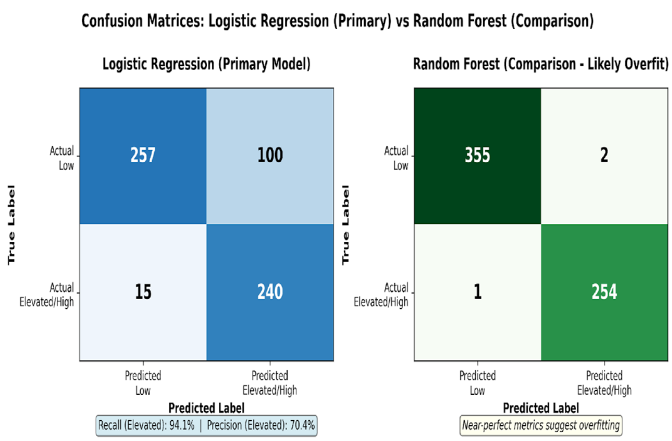

Figure 6. Confusion matrices for Logistic Regression (left) and Random Forest (right). Logistic Regression prioritizes recall, minimizing dangerous false negatives at the cost of some false positives - an appropriate trade-off for public safety applications. Random Forest shows near-zero training error but higher test-set misclassification, indicating overfitting.

5.3 Proposed Improvements

Based on the experimental results and architectural analysis, several systematic improvements have been identified to enhance the system's operational capability. The current Data Sources Layer could be significantly refined by integrating high-resolution vegetation density data from LANDFIRE and topographic exposure metrics, which would provide a more granular assessment of local fuel conditions. Additionally, the transition from retrospective analysis to proactive forecasting requires the incorporation of real-time forecast data from the NOAA National Weather Service, allowing the system to provide multi-day projections rather than just current-state assessments. While the current study prioritized interpretability via Logistic Regression, future iterations propose the exploration of Bayesian hyperparameter optimization to improve precision without compromising the high recall standards established in this study. Finally, the modular architecture is designed to be hazard-agnostic, suggesting a proposal to expand the system to manage other rapid-onset events, such as flash floods, by integrating domain-specific sensors like USGS stream gauges.

5.4 Validation

The reliability of the system was rigorously verified through temporal consistency validation to ensure that the predictive performance was not an artifact of data leakage. A specific case study analysis, visualized in Figure 7, tracked the "Distance to Nearest Fire" feature over time relative to risk thresholds. The pattern demonstrates that the system strictly adheres to chronological causality; high distances are recorded before ignition, and the value decreases only as the fire approaches. This confirms that the model's filtering logic successfully isolates future fire data from the training process, ensuring that the risk assessments reflect a realistic operational environment where only past and current information is available.

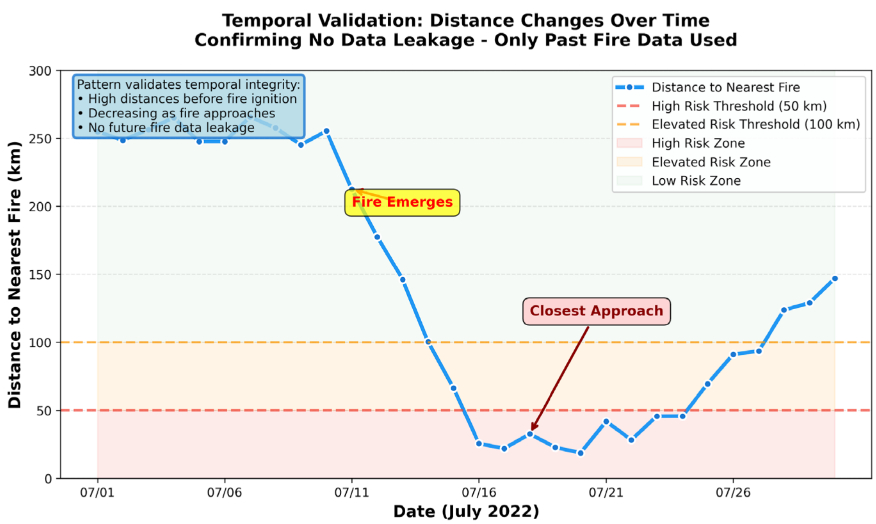

Figure 7. Temporal validation of the data leakage prevention protocol. Distance-to-nearest-fire remains high before ignition and drops only after a fire event begins, confirming that the model's filtering logic maintains strict chronological causality and the retrospective performance metrics reflect realistic operational conditions.

5.5 Discussion

The implementation of the Wildfire Readiness Assistant highlights the critical role of interpretability and human-centered design in public safety AI. Unlike "black-box" deep learning models, the use of an interpretable Logistic Regression framework allows users and emergency managers to clearly understand the physical drivers—specifically fire proximity and wind speed—behind each readiness alert. This transparency is essential for reducing cognitive load and fostering user trust during high-stress scenarios. Furthermore, the system's architecture explicitly enforces a "Human/Official Boundary," as detailed in Figure 1, which ensures that the AI functions solely as a supplementary advisory tool. By strictly deferring evacuation mandates to official authorities like CAL FIRE, the system avoids the ethical and safety risks associated with autonomous decision-making in disaster contexts. This approach successfully reframes the challenge from stochastic fire prediction to deterministic readiness management, providing a scalable template for enhancing community resilience.

5.6 Limitations

While this system demonstrates effective performance on the 2022 California wildfire season dataset (16 geographic tiles, 2,448 location-day observations), several limitations warrant acknowledgment and suggest directions for future work.

Geographic and Temporal Scope: The model was trained and validated exclusively on California data from the 2022 fire season. Generalization to other geographic regions with different fire regimes (e.g., grassland vs. forest fires) or future fire seasons under evolving climate conditions may require model recalibration or retraining. The Santa Cruz County pilot region, used during initial development for pipeline validation and proof-of-concept testing, exhibited highly separable fire-proximity patterns that enabled rapid iteration but was insufficient for rigorous model evaluation. All performance metrics reported in this study reflect the broader statewide California dataset.

Data Availability and Real-Time Constraints: The current implementation relies on retrospective fire perimeter data from CAL FIRE FRAP, which provides verified fire geometries but with inherent reporting delays. Transition to real-time operational forecasting would require integration with predictive weather models and active fire detection systems (e.g., satellite-based MODIS or VIIRS), introducing additional sources of uncertainty.

Spatial Resolution: The tile-centroid approach, where weather data is fetched from a single point representing each geographic tile, provides household-relevant regional readiness guidance. However, this aggregation may not capture microclimatic variations or topographic effects within complex terrain, potentially affecting accuracy in areas with significant elevation gradients or localized weather patterns.

User Acceptance and Behavioral Factors: The system's operational effectiveness ultimately depends on user trust, comprehension of readiness levels, and willingness to act on guidance. These human factors would require longitudinal field studies, user surveys, and deployment trials to properly assess. Additionally, the system is designed to complement—not replace—official emergency alerts and evacuation orders from authorities.

Model Interpretability vs. Performance Trade-offs: Logistic Regression was selected for its interpretability in safety-critical applications, accepting lower performance (ROC-AUC 0.92) compared to more complex models like Random Forest (ROC-AUC 0.96). This trade-off prioritizes transparency and stakeholder trust but may limit predictive ceiling.

Model performance on the High readiness tier cannot be assessed from this dataset, as no observations meeting all three simultaneous threshold conditions (proximity < 25 km, temperature > 35°C, wind speed > 30 km/h) were recorded during the 2022 fire season. This is an expected characteristic of conservatively defined rare-event categories; however, it means the three-tier classification system has only been empirically validated on two of its three classes.

6. Conclusion

This research successfully demonstrates that AI-driven classification, when implemented within a rigorous systems engineering framework, can effectively bridge the critical operational gap between institutional emergency management and household preparedness. By utilizing an interpretable Logistic Regression classifier specifically optimized for recall, the Wildfire Readiness Assistant (WiRA) achieved a 94.1% recall rate in identifying elevated-risk conditions (evaluated on the binary Elevated vs. Low classification task, as no High-risk observations were present in the 2022 dataset), a performance metric validated through strict temporal consistency testing to ensure scientific integrity.

Beyond quantitative accuracy, the study establishes a significant contribution to responsible AI design philosophy. As formalized in the system architecture, the implementation of an explicit "Human/Official Boundary" ensures that artificial intelligence serves strictly to augment human judgment rather than usurp the authority of official emergency directives. By reframing the challenge of wildfire risk from a stochastic prediction problem to a deterministic readiness management task, this work provides a scalable and transparent template for future public safety applications. As climate-driven environmental hazards continue to intensify, the synergy of authoritative data integration, machine learning, and human-centered design principles remains essential for fostering resilient communities and ensuring safe AI deployment in high-stakes environments.

References

Abatzoglou, J. T. and Williams, A. P., "Impact of Anthropogenic Climate Change on Wildfire Across Western U.S. Forests," *Proceedings of the National Academy of Sciences*, vol. 113, no. 42, pp. 11770–11775, 2016.

Ashfaq, H., "Harnessing AI for Wildfire Defense: An Approach to Predict and Mitigate Global Fire Risk," *Proceedings of the NeurIPS 2024 Workshop on Tackling Climate Change with Machine Learning*, 2024.

Berkhahn, S., Fuchs, L., and Neuweiler, I., "An Ensemble Neural Network Model for Real-Time Prediction of Urban Floods," *Journal of Hydrology*, vol. 575, pp. 743–754, 2019.

Breiman, L., "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

CAL FIRE, "Fire and Resource Assessment Program (FRAP)," California Department of Forestry and Fire Protection, 2023. Accessed Jan. 2026.

Finney, M. A., "FARSITE: Fire Area Simulator—Model Development and Evaluation," USDA Forest Service Research Paper RMRS-RP-4, 1998.

George, M. B. and Ayiku, E. O., "AI-Driven Fire Risk Indices Integrating Climate, Fuel, and Terrain for Wildfire Prediction and Management," *International Journal of Engineering Technology Research & Management*, vol. 8, no. 2, 2024.

Giglio, L., Schroeder, W., and Justice, C. O., "The Collection 6 MODIS Active Fire Detection Algorithm and Fire Products," *Remote Sensing of Environment*, vol. 178, pp. 31–41, 2016.

Iqbal, U., Perez, P., Li, W., and Barthelemy, J., "How Computer Vision Can Facilitate Flood Management: A Systematic Review," *International Journal of Disaster Risk Reduction*, vol. 53, 102030, 2021.

Jain, P., Coogan, S. C. P., Subramanian, S. G., Crowley, M., Taylor, S., and Flannigan, M. D., "A Review of Machine Learning Applications in Wildfire Science and Management," *Environmental Reviews*, vol. 28, no. 4, pp. 478–505, 2020.

Lundberg, S. M. and Lee, S.-I., "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

NASA, "Fire Information for Resource Management System (FIRMS) Documentation," NASA Earth Science Data Systems, 2024. Accessed Jan. 2026.

National Wildfire Coordinating Group, "National Fire Danger Rating System (NFDRS) Documentation and Updates," 2023. Accessed Jan. 2026.

Open-Meteo, "Historical Weather API Documentation," Open-Meteo Archive, 2024. Accessed Jan. 2026.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., and Duchesnay, É., "Scikit-Learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

Ribeiro, M. T., Singh, S., and Guestrin, C., "Why Should I Trust You? Explaining the Predictions of Any Classifier," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1135–1144, 2016.

Shneiderman, B., "Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy," *International Journal of Human–Computer Interaction*, vol. 36, no. 6, pp. 495–504, 2020.

Shukla, M., et al., "Tailoring Household Disaster Preparedness Interventions to Reduce Health Disparities: Nursing Implications from Machine Learning Importance Features," *International Journal of Environmental Research and Public Health*, vol. 21, no. 5, 521, 2024.

Van Wagner, C. E., "Development and Structure of the Canadian Forest Fire Weather Index System," Canadian Forestry Service, Forestry Technical Report 35, 1987.

Westerling, A. L., "Increasing Western U.S. Forest Wildfire Activity: Sensitivity to Changes in the Timing of Spring," *Philosophical Transactions of the Royal Society B*, vol. 371, no. 1696, 2016.

Yaras, I. E. and Onan, A., "Artificial Intelligence Assisted Decision Support System for Forest Fire Crisis Management," *Scientific Journal of Mehmet Akif Ersoy University*, vol. 8, no. 2, pp. 97–112, 2025.

Biography

Anshul Dhaas is a junior at Westmont High School in Campbell, California, USA with interests in machine learning, public safety, and engineering for community impact. His project, the Wildfire Readiness Assistant, applies interpretable machine learning to help households assess wildfire readiness using weather and fire perimeter signals. He is passionate about building responsible AI tools that support public safety decision-making and complement official emergency alerts.