WiSE 2.0: An Integrated Probabilistic Platform for Site-Specific Wildfire Egress Planning and Decision Support

Wadie Chalgham[a], Taehwan Lim[a], Carl Swindle[a], and Ali Mosleh[a]
[a] The B. John Garrick Institute for the Risk Sciences, University of California Los Angeles, USA
Emails: Wadie.chalgham@ucla.edu, Taelim@g.ucla.edu, Mosleh@g.ucla.edu

---

Abstract: Wildfires worldwide have become increasingly large and destructive in recent decades, resulting in catastrophic damage and evacuation failure. In the United States, many communities lack the necessary planning strategies to optimize for successful fire evacuation. While limited-capability technologies are emerging, no platform exists that combines fire spread, routing and traffic, and human behavior models into a fully integrated solution with realistic capability to meet the needs of stakeholders such as city planners, emergency responders, vegetation managers, and community educators. This paper presents the Wildfire Safe Egress (WiSE) 2.0 framework and its accompanying decision-support platform, which integrate fire dynamics, human decision-making, and traffic modeling within a probabilistic Bayesian Belief Network to estimate the likelihood of successful community evacuation under different scenarios. WiSE 2.0 advances earlier integrated frameworks along three axes: a high-throughput HFire-based fire spread model that is approximately 160 times faster than the FARSITE/FlamMap baseline; a microscopic SUMO-based traffic module accelerated for Monte Carlo analysis and a Census- and Social Vulnerability Index (SVI)-grounded human behavior model that samples per-vehicle pre-evacuation times from a calibrated demographic regression. The framework is implemented as a Tauri-packaged desktop client and a browser-deployed web application sharing an Axum server, with a GIS-based interface that enables stakeholders to construct site-specific evacuation scenarios across the United States, evaluate road-networks and warning-system interventions, and compare outcomes through a built-in sensitivity-analysis dashboard and a Bayesian-inference view. The framework is demonstrated through a calibration and validation case study based on the 2018 California Camp Fire. WiSE 2.0 is ready for collaboration with stakeholders for further evaluation.

Keywords: Wildfire Evacuation, Probabilistic Risk Assessment, Bayesian Belief Network, Safe Egress, HFire, SUMO, Decision Support System.

1. INTRODUCTION

Wildfires pose a significant and escalating threat to communities in Wildland-Urban Interface (WUI) areas worldwide. According to the National Interagency Fire Center, in 2021 alone there were 58,985 wildfires in the United States, burning approximately 7.1 million acres [1]. Wildfires cause property damage, asset loss, civilian casualties, and environmental destruction every year, and they rank among the most significant natural hazards in many regions — particularly California, where nine of the ten largest wildfires on record have occurred within the past seven years, and where more than two million properties were classified at high-to-extreme wildfire risk in 2021 [2]. The January 2025 Palisades fire alone, destroyed more than 6,800 structures and triggered mandatory evacuation orders that extended into Malibu within hours of ignition.

Wildfire risk management spans prevention, mitigation, and emergency response, of which the safe evacuation of WUI communities is one of the most consequential and least predictable phases. Wildfire egress planning is an active multi-disciplinary research area, but ensuring the safety of a WUI community in the case of a wildfire remains a complex task. It requires a deep understanding of wildfire dynamics and spread mechanisms, human decision-making processes under emergency conditions, and transportation system response under demand surges. Transportation engineers have largely focused on vehicular traffic modeling, concentrating on road networks and travel times during evacuation [3–5]. Such agent-based methods support comparison of evacuation strategies [6], simulation of multi-modal evacuation [7], and incorporation of new data sources such as cellular trajectories [8]. Social scientists, by contrast, have focused on human behavior and emergency decision-making, including pre-evacuation timing, warning compliance, and household-level mobilization [9–11].

Despite recent efforts to converge these aspects of wildfire evacuation [12, 13], no platform integrates all three elements — fire spread, human decision-making, and traffic — into a unified probabilistic framework with site-specific applicability and stakeholder-grade usability. As a result, planners and emergency responders typically reason about evacuation under each layer in isolation, using slow or outdated fire models, generic traffic assumptions, and qualitative behavioral judgement. Customer-discovery interviews conducted as part of the NSF I-Corps program with county-level emergency management directors, CalFIRE incident commanders, and resilience-focused vegetation managers identified the absence of such an integrated tool as the most cited operational gap.

This paper presents the Wildfire Safe Egress (WiSE) 2.0 framework and its accompanying decision-support platform. WiSE 2.0 extends the probabilistic egress framework first introduced in WiSE 1.0 [14, 24] along three technical axes and one software development axis; (1) The fire dynamics layer is driven by HFire [15], a Rothermel-based fire-spread model that produces 24-hour fire-arrival rasters approximately 160 times faster than the FARSITE/FlamMap baseline, (2) The traffic layer is reformulated as a microscopic SUMO-based simulation [16], (3) The human decision layer samples per-vehicle pre-evacuation times from a Census- and Social Vulnerability Index (SVI)-grounded demographic regression rather than a single community-level Poisson rate, and (4) These advances are packaged into a unified web and desktop platform — a Rust-based Tauri client and an Axum HTTP server that orchestrates a sandboxed Python operation pipeline — with a Leaflet-based GIS interface, a real-time sensitivity-analysis dashboard, and a Bayesian-inference view that lets stakeholders impose evidence on any node of the probabilistic model and observe updated posteriors.

The remainder of the paper is organized as follows; Section 2 reviews the WiSE 2.0 probabilistic framework and the structure of the Bayesian Belief Network, Section 3 describes the platform architecture and the scenario creation workflow, Section 4 details the three core models and the advances introduced in this work, Section 5 presents a calibration and validation case study based on the 2018 California Camp Fire, and Section 6 discusses the stakeholder deployment plan and ongoing validation efforts, and Section 7 concludes.

2. THE WISE 2.0 PROBABILISTIC FRAMEWORK

Wildfire egress planning consists of three tightly coupled layers — fire dynamics, human decision-making, and traffic — each of which affects and is affected by the others. Fire dynamics depend on available vegetation, weather, and topography. The human decision layer depends on population distribution, demographic characteristics, and the quality and timing of warning systems. Traffic depends on road-network configuration, available transport modes, and the temporal demand curve produced by the human layer. Visible fire or smoke can accelerate departure decisions; fire can damage communication infrastructure and segments of the road network; and congestion can in turn impede fire-suppression operations. Figure 1 summarizes these models and their interactions.

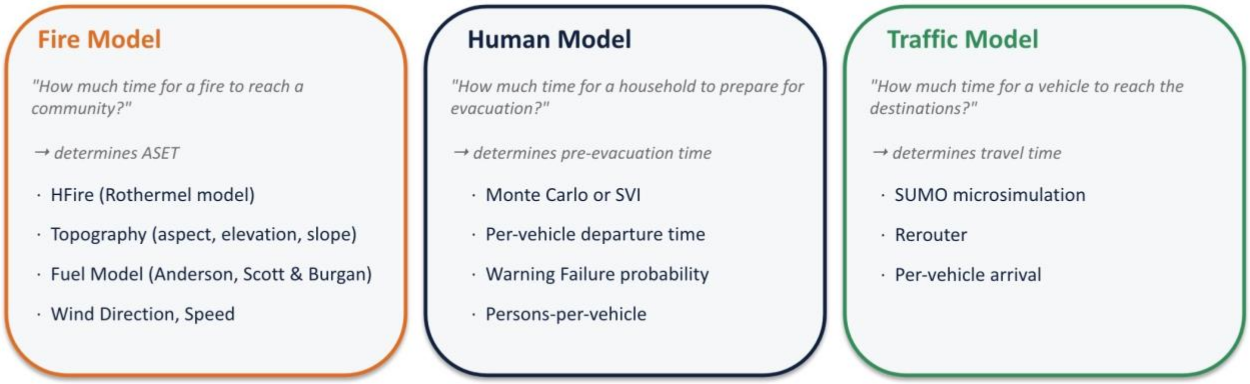

Figure 1: The main models of WiSE 2.0 Framework.

WiSE 2.0 casts the evacuation outcome as a comparison between three-time variables. The Available Safe Egress Time (ASET) is the time between fire detection and fire arrival at the community boundary; it is determined by the fire-spread model. The Required Safe Egress Time (RSET) is the time each vehicle actually needs to evacuate safely; it is the sum of a pre-evacuation component (when the household begins moving — determined by the human decision model) and a travel-time component (how long the route takes — determined by the traffic model). The Surplus Safe Egress Time (SSET) is the time at which the fire arrives at the specific road segment a vehicle is traversing; it is computed by intersecting the per-segment fire-arrival raster with each vehicle's chosen route. A vehicle evacuates safely when, for every segment of its route, RSET ≤ SSET — that is, the vehicle clears the segment before fire arrives. Probability of Safe Evacuation is then the fraction of simulated vehicles satisfying this inequality, computed via Monte Carlo simulation over the three layers and wrapped in a Bayesian Belief Network (BBN) that quantifies the uncertainty across the network.

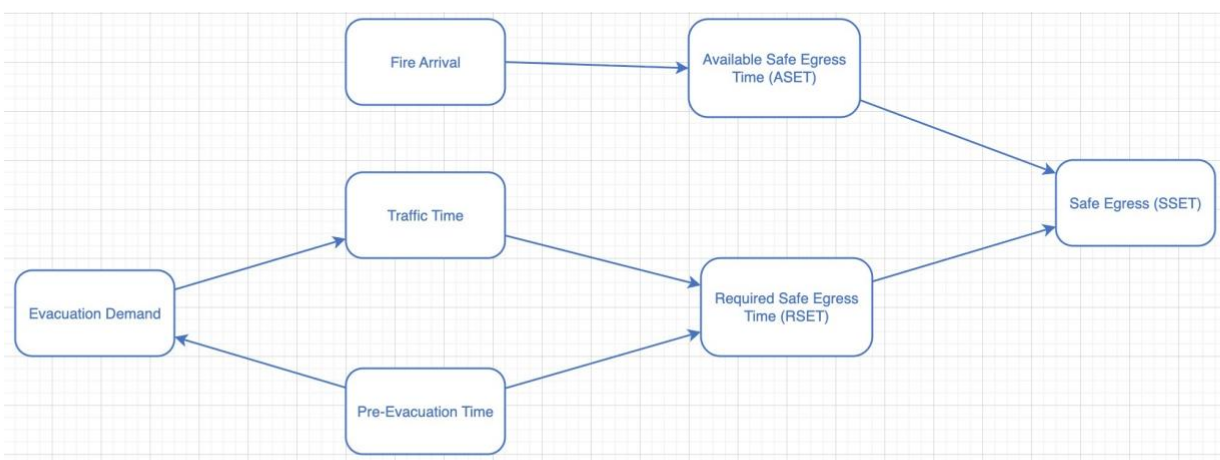

Figure 2: Bayesian Belief Network Linking the Three Models to the Probability of Safe Egress.

Counterfactual scenarios — for example, closing a specific road, adding contraflow lanes on an arterial, or changing the warning lead time — are evaluated by reseeding the underlying Monte Carlo with the altered input and rebuilding the CPTs.

3. PLATFORM ARCHITECTURE

The WiSE 2.0 framework is delivered as a stakeholder-facing decision support platform that shares a single codebase across two deployment modes: a Tauri-packaged desktop client for offline use by field responders, and a browser-hosted web application for organization-wide use. Both modes are backed by an Axum HTTP and WebSocket server written in Rust, which authenticates per-organization access, mediates a per-project disk-backed object store, and dispatches long-running computations to a sandboxed pool of Python operations. Each Python operation encapsulates a single step of the simulation pipeline (HFire fire spread, OSM extraction, SUMO trip generation, SUMO routing, traffic simulation, statistics aggregation, BBN inference) and is invoked as a subprocess via a typed registry. Long-running operations stream stdout to the client over the WebSocket job channel and report cancellation cleanly, so a user can interrupt a multi-minute SUMO run from the UI without leaving orphaned processes.

The frontend is a Leaflet-based GIS canvas with first-class support for polygon and polyline drawing, layered overlay of fire perimeters and danger zones, and a time slider for animating fire spread and vehicle trajectories. Shared geospatial datasets — OpenStreetMap road networks, U.S. Census ACS tracts and CDC SVI components, the WorldPop 1 km × 1 km gridded population [17], the Landfire 2020 fuel-and-topography rasters, the UCAR WSAP 10-hour Dead Fuel Moisture analysis [21], and the Open-Meteo hourly wind forecast [22] — are pre-staged on the server.

In the WiSE 2.0 workflow, a stakeholder first creates a project (an example is shown in Figure 3) by specifying a bounding box, a reference fire model, and the state and county. The platform then automatically extracts the local road network and demographic profile. The user then composes one or more scenarios on top of that project. A scenario is defined by five user-editable elements: (i) a community polygon, (ii) one or more shelter zones with optional safe-zone-limit polylines, (iii) optional danger-zone polygons, (iv) optional road-network alterations — closures, lane additions, and contraflow segments, each with an associated cost estimate — and (v) a warning system parameterized by mean warning time to community, warning failure probability, and persons per vehicle.

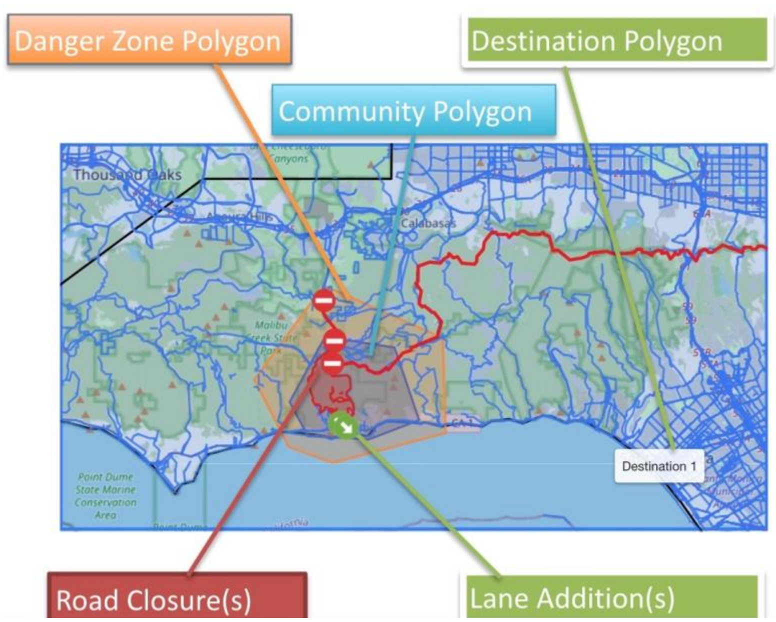

Figure 3: An example of project creation in WiSE 2.0

Each scenario also selects an Available Safe Egress Time (ASET) mode that determines how fire arrival is computed: (a) Fire Arrival mode, which derives per-segment SSET from the HFire fire-arrival raster of a selected fire model; and (b) Tenability Time mode, which parameterizes SSET as a triangular distribution between user-supplied minimum, mean, and maximum tenability times when no fire model is available.

Once a scenario is configured, the platform runs the multi-stage processing pipeline: OSM-to-SUMO network conversion, MC or SVI demographic sampling, trip generation and shortest-path routing via duarouter, road-closure post-processing through SUMO's rerouter mechanism, the SUMO microsimulation itself, and finally the per-vehicle statistics aggregation. The aggregator produces a simulation record containing, for each vehicle, departure time, travel time, RSET, ASET, the per-route SSET, and a binary safe-evacuation flag. From these records the platform derives the headline outcome metrics — Probability of Safe Evacuation, Clearance Time (the maximum RSET over safely-evacuated vehicles), Mean and P95 RSET, Mean and P5 SSET–RSET margin, and Close Calls (the percentage of vehicles with margin under five minutes) — and surfaces them in a Results page, a Sensitivity Analysis dashboard for scenario-versus-scenario comparison, and a Bayesian-Network tab that exposes the BBN structure with live posterior updates under user-imposed evidence.

4. CORE MODEL ADVANCES

4.1 Fire Dynamics: Fire Spread Modeling with HFire

The fire-spread layer determines both ASET (the community-level fire-arrival time) and SSET (per-segment fire-arrival times along candidate evacuation routes). WiSE 2.0's fire dynamics layer is driven by HFire [15], a Rothermel-based surface fire-spread model. HFire ingests fuel-model rasters (Anderson 13 or Scott & Burgan 40), Landfire 2020 topography rasters (slope, aspect, elevation, canopy cover), dead and live fuel moistures (DFM, LFM), and hourly wind speed and direction at a user-specified ignition point and start time. Dead fuel moisture is fetched at hourly resolution from the UCAR WSAP 10-hour analysis [21] for past dates; for forecast dates, the platform falls back to the most recently available analysis hour and propagates that value forward — a documented trade-off that is communicated to the user. Live fuel moisture is read from a pre-computed monthly climatology of 2016–2020 raster snapshots indexed by month, which is independent of the calendar year and therefore usable for both historical and forecast scenarios.

Hourly wind is fetched from Open-Meteo [22] with a forecast horizon of up to 16 days; for past dates, the platform automatically switches to the ERA5 archive endpoint. Wind is sampled at every cell of an upscaled aspect grid (typically 10 × 10 over the bounding box at a coarsen factor of 100), giving HFire a true spatial wind field rather than a single bbox-average vector. The HFire run produces per-hour fire perimeters in EPSG:4326 GeoJSON and a continuous 24-hour fire-arrival raster. From this output, the platform computes per-segment SSET by intersecting each SUMO edge with the temporal arrival raster. HFire executes approximately 160 times faster than FARSITE/FlamMap [23], which is what makes the Monte Carlo workflow over hundreds of fire scenarios tractable on commodity hardware.

4.1.1 Spatial Wind and Live Inputs

A common shortcoming of operational fire-spread models is the use of a single bbox-average wind vector. WiSE instead samples wind at every cell of an upscaled aspect grid (typically 10 × 10 over the bounding box at a Landfire coarsen factor of 100), giving HFire a true spatial wind field. Wind is fetched from Open-Meteo with a forecast horizon of up to 16 days; for past dates, the platform automatically switches to the ERA5 archive endpoint. The wind table for each scenario is editable from the prepare screen so a planner can override the forecast with locally-observed values, for example after a CalFIRE remote-automated-weather-station (RAWS) reading. Figure 4 shows the live spatial-wind editor in the WiSE platform; each row corresponds to one upscaled grid cell and one hour of the simulation.

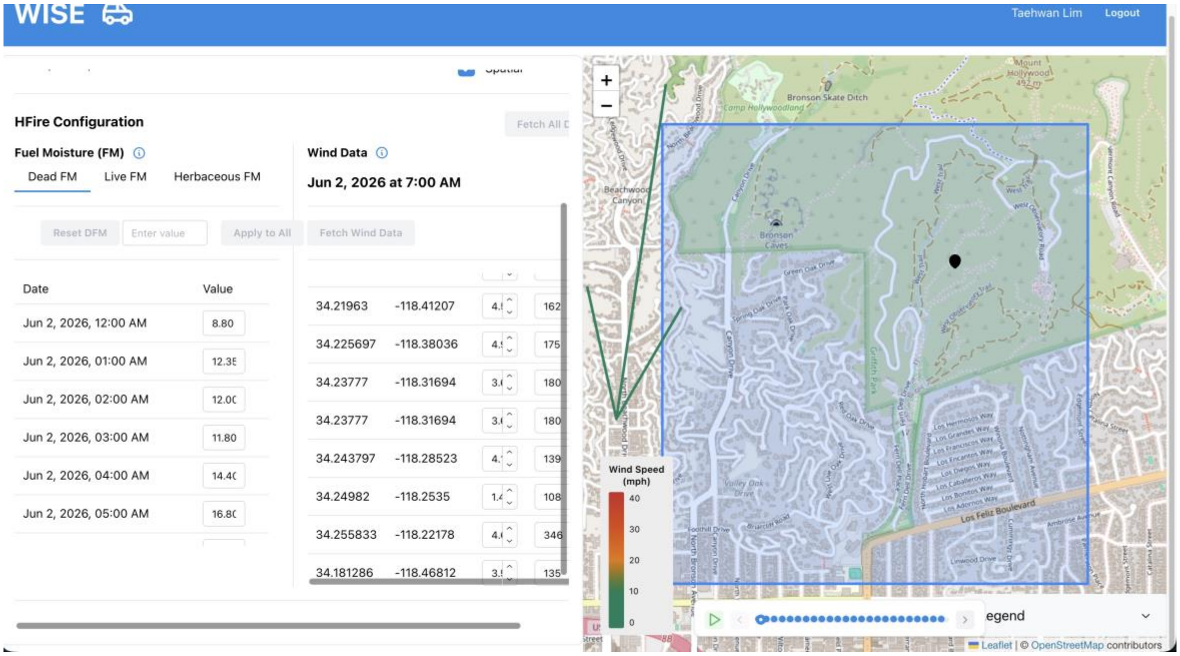

Figure 4: Spatial Wind Editor in the WiSE 2.0 HFire Configuration View.

Dead fuel moisture is fetched at hourly resolution from the UCAR WSAP 10-hour analysis [21] for past dates. WSAP is an analysis product, not a forecast, so for forecast dates the platform falls back to the most recently available analysis hour and propagates that value forward — a documented trade-off that is communicated to the user in the wind-data panel. Live fuel moisture is read from a pre-computed monthly climatology of 2016–2020 raster snapshots indexed by month and is therefore independent of calendar year, making it usable for both historical and forecast scenarios.

4.1.2 From Perimeter Raster to ASET and SSET

Each completed HFire run produces a time-indexed set of fire perimeters; the latest perimeter is persisted as the canonical footprint of the model and the per-hour perimeters are written to a timestamped GeoJSON directory. To compute ASET for a scenario, the platform finds the earliest hour at which the fire perimeter intersects the user-drawn community polygon. To compute SSET for each road segment, the platform rasterizes the time-indexed perimeters at 30 m resolution and reads the fire-arrival time for every SUMO edge that the routing layer subsequently emits as a candidate route. Figure 5 shows the per-fire-model dashboard with the active perimeter overlay and the daily summary panel; this is the user-facing surface where a planner reviews and tunes the fire input before composing scenarios.

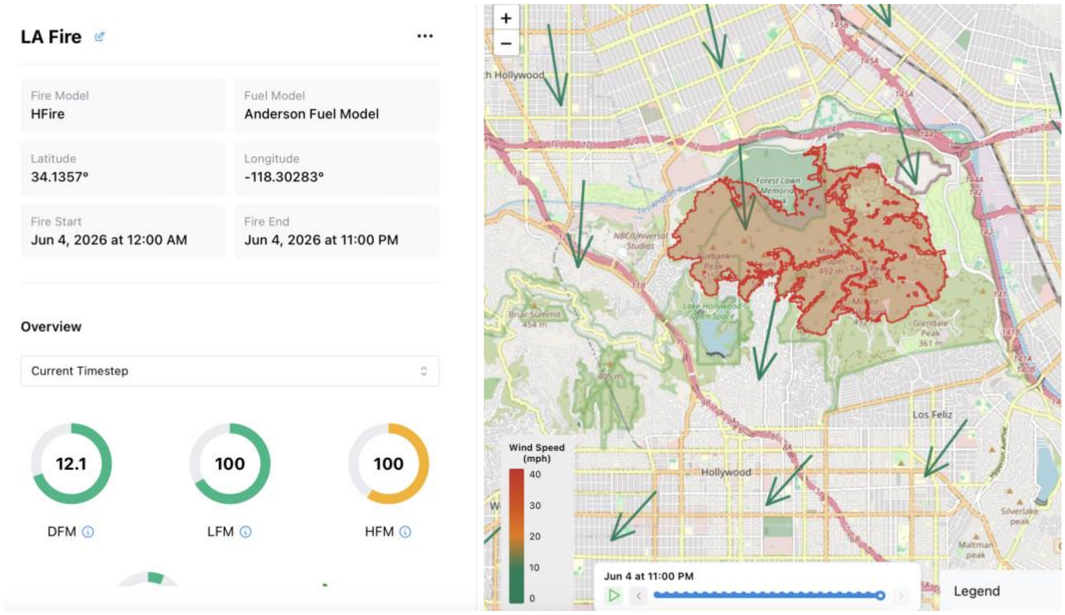

Figure 5: An example of WiSE 2.0 Fire Model Overview with HFire Output and Daily Fuel-Moisture Summary.

4.2 Human Decision-Making Modeling

The human decision layer models when each community member begins to evacuate — the departure-time component of RSET. The community polygon is intersected with the WorldPop 1 km × 1 km gridded population [17] and then disaggregated into vehicles using a user-configurable persons-per-vehicle ratio. WiSE 2.0 supports two demographic behavior models, both invoked through a shared Python op interface. Figure 6 shows an example of WiSE 2.0 Human model demographics data, that are autofilled from census data, but the user has the option to edit them.

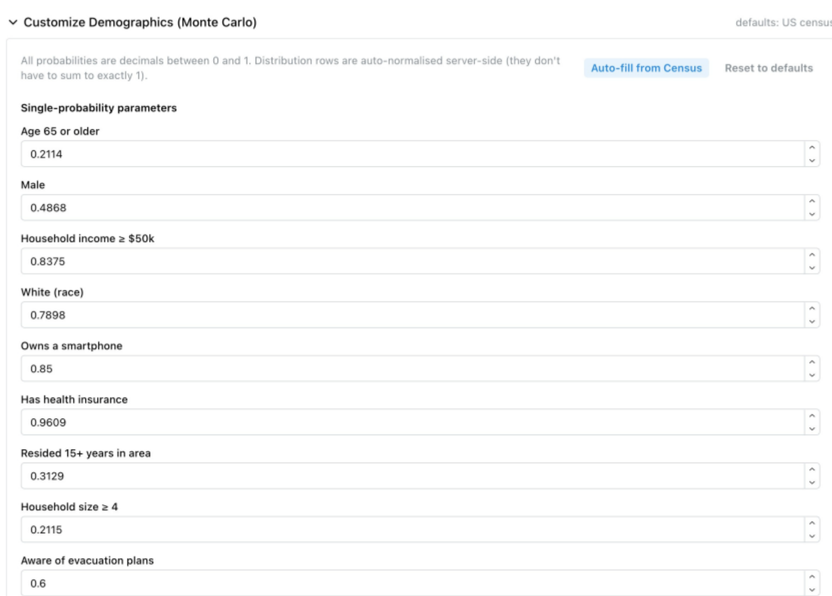

Figure 6: An example of WiSE 2.0 Human model demographics data.

4.2.1 Two Demographic Behavior Models

WiSE supports two interchangeable demographic behavior models, both invoked through a shared Python operation. The Monte Carlo (MC) model samples per-vehicle pre-evacuation times from a behavioral regression conditioned on community-level demographic priors. The default prior vector mirrors U.S. Census national averages and includes proportion of residents aged 65 or older, proportion male, proportion in households with income above $50,000, proportion white, proportion smartphone-owning, proportion insured, proportion resident for 15 or more years, proportion in households of four or more, and proportion familiar with the local evacuation plan; together with categorical distributions over the number of transportation modes available, the number of evacuations experienced, the educational attainment level, and the primary mode by which the household learns of the warning. Each of these is editable so a planner can match local survey data when available.
The Social Vulnerability Index (SVI) model derives the demographic vector directly from CDC SVI tract-level components [19] within the community polygon, weighted by population, and applies the same regression to produce a per-vehicle pre-evacuation distribution. The SVI model is well-suited to communities with high within-tract heterogeneity, where a single community-level prior would mask vulnerable sub-populations.

4.2.2 Warning System Layer

Both behavior models feed into a warning-system layer characterized by a warning failure probability — the fraction of the community that never receives the official evacuation notice and must rely on secondary cues such as visible fire, social contagion, or notification from neighbors — together with a mean warning lead time and a triangular distribution over the minimum, mean, and maximum tenability times. The product of the two layers is a per-vehicle departure-time distribution: an exponentially-tailed function for vehicles that received the official warning, plus a heavier-tailed function for vehicles that did not. This distribution is then sampled to seed the traffic simulation's per-vehicle origin-departure times.

4.3 Traffic Modeling: SUMO

The traffic layer determines the travel-time component of RSET. Microscopic traffic is simulated using SUMO [16]. The OpenStreetMap road network within the project bounding box is extracted via osmium and converted to SUMO's .net.xml format using netconvert. Trips are generated by sampling per-vehicle origins uniformly inside the community polygon and snapping each to the nearest network node, and assigning destinations based on a probabilistic distribution over the configured shelter zones. duarouter computes shortest paths over edge attributes that include road type, lane count, and segment length. The router honors user-defined road-network alterations — full edge closures, lane additions, and contraflow segments — and a post-processing step truncates each route at the first edge that enters a shelter polygon so that intra-shelter movement is not double-counted in RSET.

4.3.1 Road Closures and Time-Stepped Fire Closures

Road closures are implemented through SUMO's rerouter mechanism. The platform generates a rerouter additional file containing static closures derived from the user's drawn closure polylines (each polyline is matched to its underlying SUMO edge by a midpoint-distance check with a 20 m buffer, which empirically matches the OSM-to-SUMO offset budget) and time-stepped closures derived from the HFire fire perimeter (each per-hour fire polygon is intersected with the network to identify edges that should close as the fire advances). The two closure sets are written into the same <additional> file so SUMO's router can re-route vehicles around both static and dynamic obstructions in a single pass. Figure 7 shows a Camp Fire scenario with manual Honey Run closures (red prohibition markers) and the fire-derived danger zone (orange overlay).

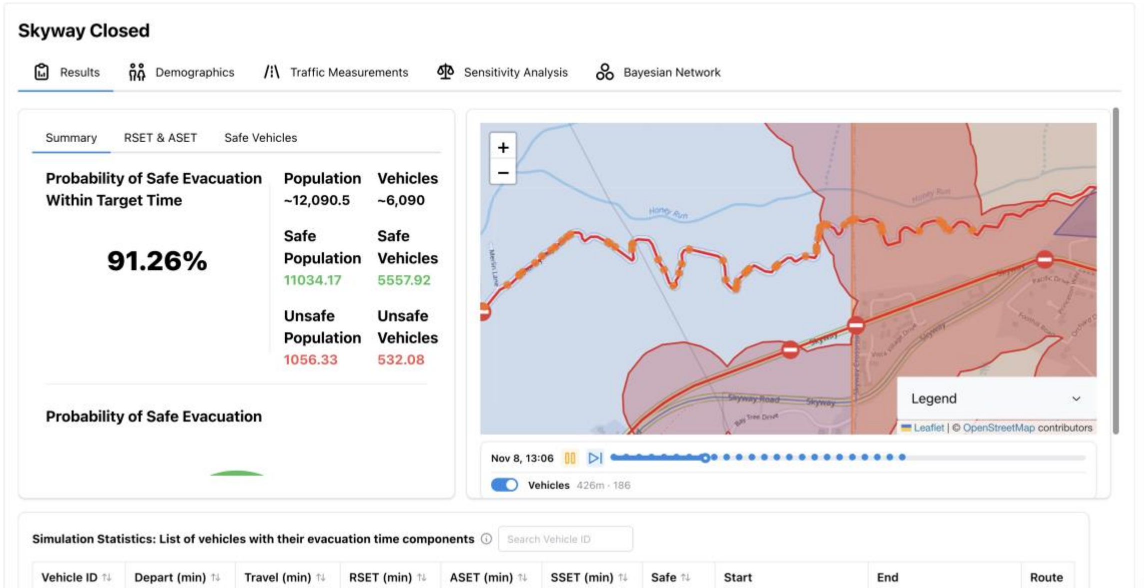

Figure 7: WiSE Scenario with Manual Road Closures and Time-Stepped Fire-Derived Closures.

4.4 Integration: From Three Models to a Probability of Safe Egress

The three models above feed five random variables in a Bayesian Belief Network. The Departure-time node is sampled from the human decision model. The Travel-time node is sampled from the traffic model conditioned on the realized departure time (and therefore on instantaneous network congestion). The Required Safe Egress Time node deterministically sums departure and travel times via an overlap convolution. The Safe Egress Time node samples a per-segment fire-arrival time from the HFire output, conditioned on the route each vehicle takes. The terminal Safe node returns a binary indicator that, for every segment on each vehicle's route, RSET does not exceed SSET — equivalently that ASET, the earliest community-level fire arrival, is itself bounded above by the route-specific SSET.

Each of the five conditional probability tables (CPTs) is estimated empirically from the underlying Monte Carlo simulation. CPT_0 is the marginal P(depart), built from 20 quantile-derived bins over the per-vehicle departure-time samples. CPT_1 is P(travel | depart), built from the joint distribution of departure-and-travel times across all simulated vehicles. CPT_2 is P(RSET | depart, travel) and is closed-form: the overlap of the depart and travel interval distributions weighted by their lengths. CPT_3 is P(SSET | RSET) and is estimated from the joint distribution of route-segment fire-arrival times against RSET. CPT_4 is the terminal P(safe | SSET) — a binary mapping implemented as a categorical with explicit support {0, 1} so that homogeneous outcomes such as 100% safe vehicles yield a well-formed 2 × n_SSET table rather than a degenerate 1 × n_SSET that would break downstream inference.

Inference is performed by the mass_BBN engine [20] using exact message passing. The small node degree of the network allows the full posterior to be computed in under one second, which makes the BBN suitable for live use in the stakeholder dashboard: a user can impose evidence on any node and observe the updated posterior for Safe Egress directly in the interface, without having to rerun the underlying Monte Carlo. Figure 8 shows the live Bayesian-Network view, in which Probability of Safe Evacuation is updated in real time as the user adjusts the Departure node's distribution to model delayed-mobilization scenarios.

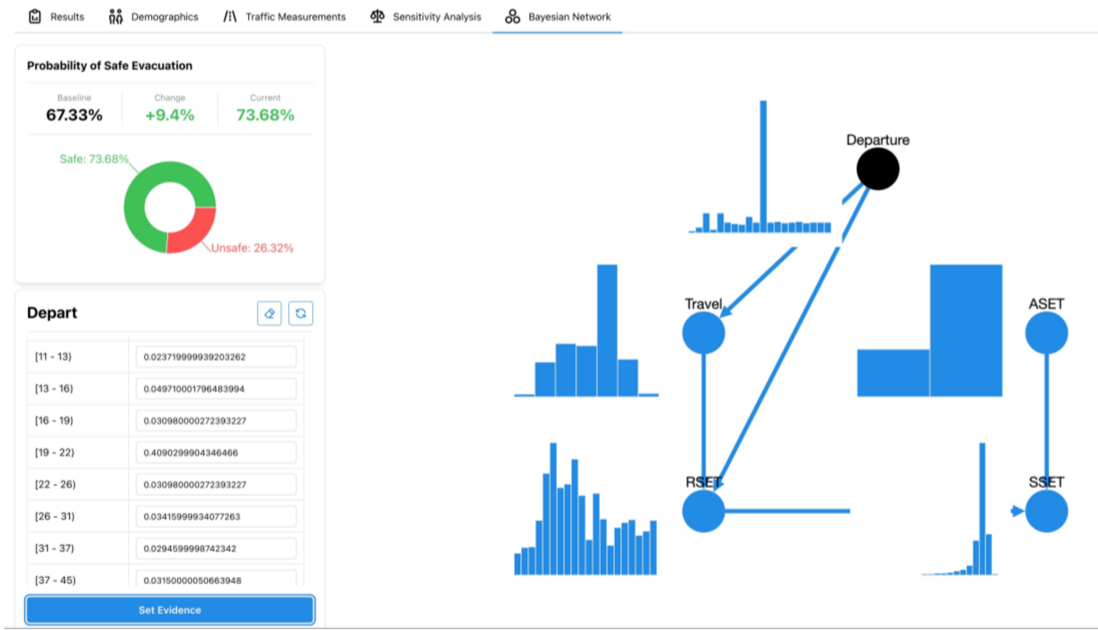

Figure 8: Live Bayesian-Network Inference View in the WiSE 2.0 Dashboard.

5. CASE STUDY: 2018 CALIFORNIA CAMP FIRE

To demonstrate WiSE end-to-end and to calibrate its parameters, the framework is applied to the 2018 California Camp Fire, which destroyed the town of Paradise and caused at least 85 civilian fatalities [18]. The project bounding box covers Paradise and the surrounding Butte County WUI. The fire is simulated in HFire using the recorded November 8, 2018 ignition point at Pulga and the corresponding hourly wind record. The warning system is configured with a mean lead time of one hour before the fire entered Paradise — consistent with the NIST timeline [18] — reflecting that the formal evacuation notification reached the community late. Under this baseline, WiSE estimates a Probability of Safe Evacuation of approximately 16 percent at the moment the fire crosses the community boundary, in close agreement with responders' first-hand estimates of 20 percent of residents evacuated at fire arrival.

Counter-factual scenarios reveal subtler structure. Increasing the effective warning lead time (and therefore ASET) from 1 to 4 hours raises the predicted safe-evacuation probability above 88 percent; extending ASET to 6.5 hours yields near-complete evacuation. Closing the Honey Run Road corridor — historically the most fatal egress route during the actual Camp Fire because it traces a steep canyon directly under the fire-spread path — improves the Probability of Safe Evacuation by approximately 1.5 percentage points relative to the no-closure baseline, because the displaced demand is routed onto Skyway, which lies further from the fire front and therefore has a higher per-segment SSET.

6. STAKEHOLDER VALIDATION AND DEPLOYMENT

WiSE 2.0 is currently ready for close collaboration with end-user stakeholders. Customer-discovery interviews conducted as part of the NSF I-Corps Hub West Region program (Award IIP-2048703) with county-level emergency management directors, CalFIRE incident commanders, and resilience-focused vegetation managers have identified three primary use cases: (i) evacuation-plan design, in which the user evaluates proposed road-network or warning-system changes against the baseline; (ii) operational tabletop exercises, in which responders rehearse evacuation under multiple fire-spread realizations; and (iii) public education and consequence communication, in which the platform's intuitive map-based output is used to explain risk to community members.

Common stakeholder requirements identified across these segments are: a low-friction scenario-creation workflow, defensible underlying methodology, and support for road-network alterations including contraflow, and rapid what-if comparison. The sensitivity-analysis dashboard and the Bayesian-Network view (Section 3) are direct responses to these requirements. Common concerns identified are: model fidelity under demographic edge cases (e.g., communities with high mobility-impaired populations), and the dependence of HFire output on accurate fuel-and-moisture rasters. The first concern is addressed by exposing all MC behavioral parameters in the prepare screen so a planner can match local survey data. The second is addressed by integrating the UCAR WSAP DFM analysis and a Landfire fuel snapshot directly into the workflow rather than requiring the user to source these inputs separately.

7. CONCLUSIONS

WiSE 2.0 integrates fire dynamics, human decision-making, and traffic modeling into a single probabilistic platform for wildfire egress planning. The platform addresses a clear gap in the state of the art: no prior framework combines these three components with site-specific applicability across the continental United States, stakeholder-grade usability through both desktop and web deployments, and the computational throughput needed for live Monte Carlo analysis. The three model advances introduced in this work — an HFire-based fire-spread layer that runs approximately 160 times faster than FARSITE, a SUMO traffic simulation that combines microscopic realism with sub-second counter-factual evaluation, and a Census- and SVI-grounded human behavior model — together make scenario-driven planning tractable on commodity hardware. The 2018 Camp Fire case study demonstrates that WiSE 2.0 reproduces responders' first-hand observations under historical conditions and that it quantifies the safe-egress benefit of warning-system and road-network interventions. The platform is ready for collaboration with stakeholders for further evaluation, and ongoing work will expand the validation case studies and integrate live data feeds in preparation for operational use.

Acknowledgements

This material is based upon work supported in part by the National Science Foundation through the Regional I-Corps Training Program (I-Corps Hub West Region, Award No. IIP-2048703). The authors want to thank the NSF I-Corps program, instructors, as well as all incident commanders, emergency managers, and fire marshals for sharing their knowledge and experience on wildfire evacuation and feedback on platform priorities.

REFERENCES

[1] National Interagency Fire Center, "Statistics: Total Wildfires and Acres (1983–2021)," [Online]. Available: https://www.nifc.gov/fire-information/statistics.
[2] Insurance Information Institute, "Facts and Statistics: Wildfires," (2022).
[3] H. Mahmassani, "Dynamic network traffic assignment and simulation methodology for advanced system management applications," Networks and Spatial Economics, vol. 1, pp. 267–292, (2001).
[4] T. Cova and J. Johnson, "A network flow model for lane-based evacuation routing," Transportation Research Part A, vol. 37, pp. 579–604, (2003).
[5] E. Kuligowski, R. Peacock and B. Hoskins, "A Review of Building Evacuation Models," NIST Technical Note 1680, (2010).

[6] S. Wong, S. Shaheen and J. Walker, "Wildfire Evacuation Strategies and Behavior in the Northern California Wildfires," Transportation Research Record, vol. 2675, no. 8, pp. 451–468, (2021).
[7] S. Pel, M. Bliemer and S. Hoogendoorn, "A review of travel behavior modeling in dynamic traffic simulation models for evacuation planning," Transportation, vol. 39, pp. 97–123, (2012).
[8] X. Li, P. Murray-Tuite and J. Wernstedt, "Cellular data for wildfire evacuation analysis: data quality and applications," Transportation Research Part C, vol. 132, (2021).
[9] T. Cova, P. Dennison, T. Kim and M. Moritz, "Setting wildfire evacuation trigger points using fire spread modeling and GIS," Transactions in GIS, vol. 9, no. 4, pp. 603–617, (2005).
[10] E. Kuligowski, "Evacuation decision-making and behavior in wildfires: Past research, current challenges and a future research agenda," Fire Safety Journal, vol. 120, (2021).
[11] S. Grajdura, X. Qian and D. Niemeier, "Awareness, departure, and preparation time in no-notice wildfire evacuations," Safety Science, vol. 139, (2021).
[12] E. Ronchi and S. Gwynne, "Computational Evacuation Modeling in Wildfires," Encyclopedia of Wildfires and Wildland-Urban Interface (WUI) Fires, Springer, no. 121, (2019).
[13] E. Ronchi, S. M. Gwynne, G. Rein, P. Intini and R. Wadhwani, "An open multi-physics framework for modelling wildland-urban interface fire evacuations," Safety Science, vol. 118, pp. 868–880, (2019).
[14] M. Pishahang, A. Ruiz-Tagle, E. Lopez Droguett, M. Ramos and A. Mosleh, "WISE: A Probabilistic Wildfire Safe Egress Planning Framework and Software Platform," Proceedings of the Probabilistic Safety Assessment and Management Conference PSAM 16, (2022).
[15] M. A. Finney, "HFire: a high-throughput surface fire-spread model based on the Rothermel formulation," USDA Forest Service Technical Report, (2007).
[16] P. A. Lopez et al., "Microscopic Traffic Simulation using SUMO," in Proceedings of the 21st IEEE International Conference on Intelligent Transportation Systems, pp. 2575–2582, (2018).
[17] A. Sorichetta, G. Hornby and F. Stevens, "High-resolution gridded population datasets for Latin America and the Caribbean in 2010, 2015, and 2020," Scientific Data, vol. 2, p. 150045, (2015).
[18] A. Maranghides et al., "A Case Study of the Camp Fire — Fire Progression Timeline," NIST Technical Note 2135, U.S. Department of Commerce, National Institute of Standards and Technology, (2021).
[19] Centers for Disease Control and Prevention / Agency for Toxic Substances and Disease Registry, "CDC Social Vulnerability Index (SVI)," [Online]. Available: https://www.atsdr.cdc.gov/placeandhealth/svi/index.html.
[20] A. Mosleh and the Garrick Institute, "mass_bbn: Exact inference for compact Bayesian belief networks," Software Package, UCLA, (2024).
[21] University Corporation for Atmospheric Research (UCAR), "WSAP 10-hour Dead Fuel Moisture Analysis," THREDDS Data Server, [Online]. Available: https://tds.rap.ucar.edu/thredds/.
[22] Open-Meteo, "Hourly Weather Forecast and ERA5 Archive APIs," [Online]. Available: https://open-meteo.com/.
[23] M. A. Finney, "FlamMap: A wildland fire behavior mapping and analysis program," U.S. Department of Agriculture Forest Service, Rocky Mountain Research Station, (2006).
[24] M. Pishahang, "Wildfire Safe Evacuation Planning and Management," Ph.D. dissertation, University of California, Los Angeles, (2025).