Minerva University

Wildfire Spread on a Real Deforestation Frontier: A Cellular Automaton Model for Brazil

CS166: Modeling and Analysis of Complex Systems

Prof. Tambasco

December 19, 2025

1 Summary

Wildfire risk in the Amazon frontier is driven by deforestation and climate extremes. Recent monitoring results show that Pará accounts for the largest share of deforestation in the Legal Amazon (INPE, 2024). Climate trends also matter. 2023 set a global temperature record and 2024 surpassed it, which increases heat stress and can raise fire risk during dry periods (Copernicus Climate Change Service, 2025). Our cellular-automaton simulations (using real 2000-tree-cover data) show that connectivity of forest patches dominates fire spread: narrow corridors of dense forest light up under ignition while adjacent clearings act as natural firebreaks. This implies that targeted interventions (e.g. strategically placed fuel breaks or buffer zones) and rapid suppression can greatly reduce burned area. Rapid-response fire crews (once a fire is detected) are especially effective in our model at keeping fire sizes small.

We recommend a hybrid strategy combining early detection with proactive fuel treatments along key forest corridors. These findings underscore the need to focus fire management on the heavily altered Amazon deforestation frontier, where fragmented forests create dangerous networks of flammable fuel.

2 Introduction

The state of Pará lies at the heart of the Amazon deforestation frontier. It has one of the most fragmented forest landscapes in Amazonia, with dense forest patches interlaced by cleared lands. Pará also leads the Legal Amazon in both deforestation and fire alerts. Over the past two decades Pará's cleared area has expanded markedly, especially in the southern part of the state, creating extensive forest edges and corridors ripe for fire spread. Large-scale fires have surged: for example, a recent report notes that over 17 million hectares (44.2 million acres) burned in Brazil's Amazon in 2024, a 66% jump from 2023. These alarming trends are driven by both human land-clearing and climate factors.

Western Amazon regions (including Pará) saw record warmth and drought in 2023–2024, making forests unusually dry and flammable. In short, the combination of rapid deforestation (creating the "arc of deforestation" across southern Pará) and climate-driven drought has pushed the region toward a critical fire-prone state. We therefore focus our modeling on a representative 3×3 km area

of Pará with real tree-cover data, to understand local fire behavior and inform mitigation.

3 Model Design

Our spread model is a stochastic cellular automaton on a 100×100 grid (3×3 km with 30 m resolution) from the Hansen Global Forest Change tree-cover map (year 2000) (See Figure 1. Each cell is one of three states: TREE (unburned fuel), BURNING, or BURNED. Time advances in discrete steps: a burning cell becomes burned (consumes its fuel), and it attempts to ignite its neighbors. We use an 8-neighbor (Moore) adjacency, allowing fire to spread in all directions (including diagonals).

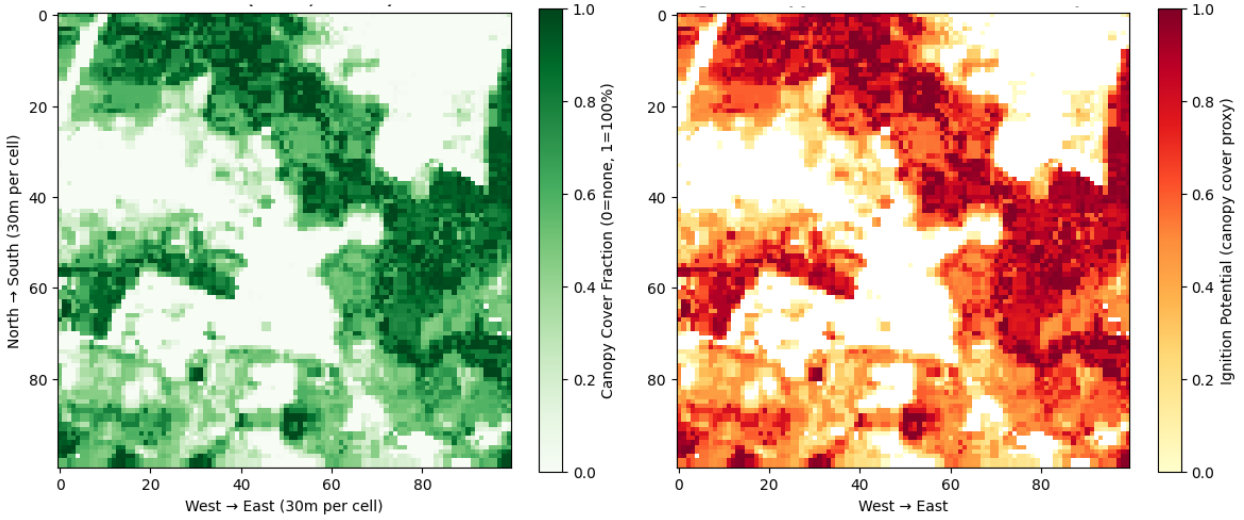

Figure 1: Tree canopy cover and ignition potential for the selected 3×3 km region in Pará. Left panel shows normalized tree cover from the Hansen Global Forest Change dataset (year 2000). Right panel shows the resulting ignition probability derived from the logistic density rule. Dense forest corridors correspond to high ignition potential, while cleared areas act as natural firebreaks.

Fire-spread is probabilistic and density-dependent. The base ignition probability for a given tree-cell depends on its local tree cover (density). Specifically, we use a logistic (sigmoid) function of density d so that cells with ¡50% cover ignite very rarely, while dense cells ignite almost certainly (See C. For example, a typical formulation is $p_{\text{base}}(d) = \frac{1}{1+\exp[-k\,(d-d_0)]}$.

k is steepness (e.g. 10) and $d_0 \approx 0.5$ (50% cover). This "S-curve" form creates a threshold-like effect: sparse vegetation is mostly immune, whereas densely wooded cells burn easily (Ghosh et al., 2024 (This approach mirrors other probabilistic fire CA models in literature, which use sigmoid

ignition functions).

Wind is incorporated as a directional multiplier on ignition probability. We fix a wind vector (direction and strength) for a simulation. When attempting to spread from a burning cell to a neighbor, we compute the angle between the neighbor direction and the wind. Downwind spread (small angle) multiplies p by up to $(1 + 1.5w)$, boosting ignition; upwind spread (angle ≈ 180) multiplies by a small factor (roughly zero), nearly halting spread. In practice this elongates the burn shape along the wind direction. All other factors (moisture, understory, spotting) are abstracted into these rules, so that tree density and wind are the main drivers.

Full technical details and formulas are given in the Appendix. In short, our percolation-based CA mimics a standard forest-fire model: fuel maps from real data, 8-way spread, density-threshold ignition, and wind-dependent transition rules. (For comparison, see the NetLogo "Fire" model or similar CA formulations). We intentionally keep the model simple to highlight how connectivity affects large-scale fire risk in this fragmented setting.

4 Parameters and Controls

The behavior of our model is governed by three primary parameters that control the physics of the spread and the logic of the interventions.

The first parameter is the sigmoid midpoint which determines the critical density required for fire to sustain itself. We set this value to 0.5 to represent a 50% canopy cover threshold. This creates a realistic tipping point where vegetation density below this level resists ignition while anything above it burns readily.

The second parameter is wind strength and it controls the directional bias of the spread. Increasing this value forces the fire to grow asymmetrically and produces elongated or elliptical fire scars similar to those driven by strong gusts in the real world.

The final parameter is the suppression threshold used specifically for the rapid response strategy. This integer represents the specific count of burned cells that triggers an intervention. It effectively models the total delay between a fire starting and the arrival of a suppression team to extinguish it.

5 Assumptions and Critical Limitations

While the model captures the topology of the landscape, we acknowledge two significant limitations that affect the quantitative accuracy of our results:

5.1 Geometric Distortion (The Diagonal Shortcut)

The model uses an unweighted Moore neighborhood, meaning a fire can spread to a diagonal neighbor with the same probability as a cardinal (North/South/East/West) neighbor. In a grid, the center-to-center distance to a diagonal cell is $\sqrt{2} \approx 1.41$, while cardinal cells are distance 1.

By treating these distances as equal, the model effectively allows fire to travel $\approx 41\%$ faster in diagonal directions than is physically realistic (Sayama, 2015). This may overestimate the rate of spread in complex, winding corridors.

5.2 Idealized Suppression

In the Rapid Response strategy, the simulation halts the fire immediately once the burned count reaches the threshold $T = 20$.

This assumes 100% containment success with zero delay once the team arrives. In reality, suppression is stochastic; containment lines can be breached, and mobilization in the dense Amazon understory takes significant time. Therefore, our Rapid Response results likely represent a "best-case scenario."

6 Results and Analysis

6.1 Wind Direction Effects

We validated the wind mechanics by running repeated simulations under different wind directions and wind strengths. Figure 2 shows that the fire footprint changes shape as the wind vector changes. With weak wind, spread is closer to radial, since each neighbor has a similar chance to ignite. With strong wind, ignition is more likely downwind and less likely upwind, so the burned area stretches in the wind direction and narrows on the back side.

This matters because the wind term changes which parts of the landscape act as corridors.

Downwind connections between dense patches become more important, and the fire can move through narrow routes that would be less reliable with no wind. At the same time, the same ignition point can produce very different outcomes depending on the wind direction, since the fire may be pushed into a dense cluster or into a clearing. This is why we treat wind as a key control for the anisotropy of spread in this model.

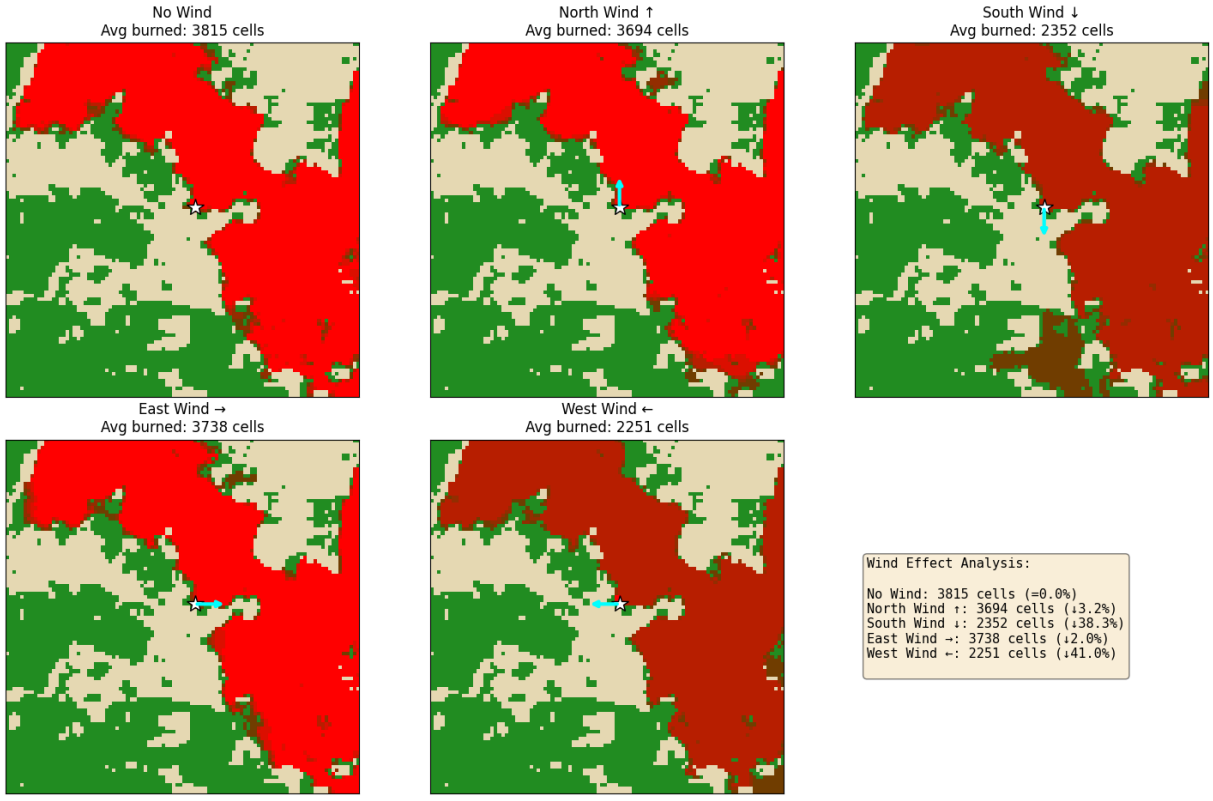

Figure 2: Effect of wind direction on wildfire spread from the same ignition point. Each panel shows the final burned area under a fixed wind direction, including no wind, north, south, east, and west winds. Strong winds bias spread downwind, producing elongated burn scars and changing which forest corridors carry fire. Reported values show average burned cells across repeated runs.

6.2 Risk Mapping and Centrality

We performed a Monte Carlo simulation ($N = 500$ trials) to generate a Risk Map. This map acts as a proxy for Betweenness Centrality in the landscape network. High-risk cells (shown in red in Figure 3) are not merely those with high density; they are the "bottlenecks" or "bridges" that connect large clusters of fuel. A fire igniting in these central nodes has the highest potential to percolate across the entire map.

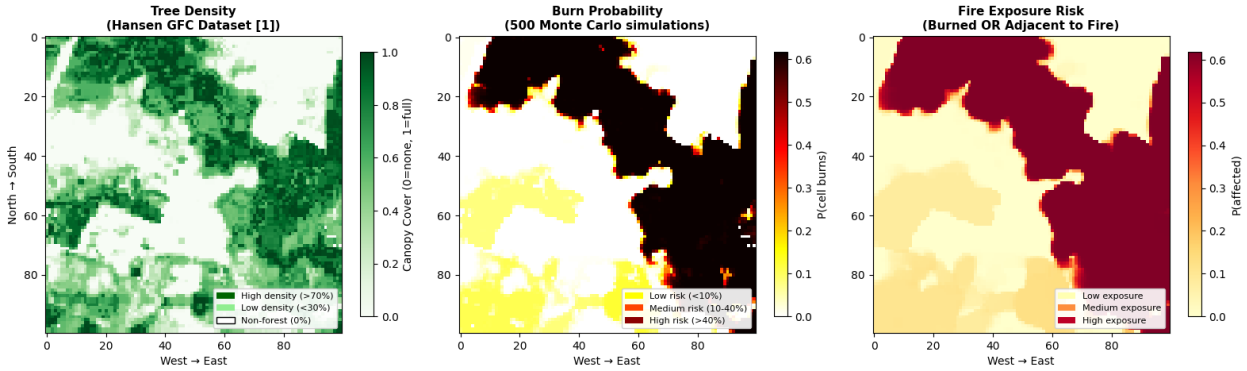

Figure 3: Risk mapping from Monte Carlo simulations. Left panel shows tree density. Middle panel shows the burn probability for each cell across 500 simulations. Right panel shows fire exposure risk, defined as cells that burn or are adjacent to burned cells. High risk regions emerge along dense corridors that connect larger forest patches.

6.3 Percolation Threshold

By scaling the global tree density, we observed a phase transition in fire behavior (Figure 4). The "spanning probability" (likelihood of fire crossing from West to East) follows a classic S-curve. The Pará landscape sits near this critical threshold, implying that small changes in connectivity (reforestation or clearing) could drastically alter fire regimes.

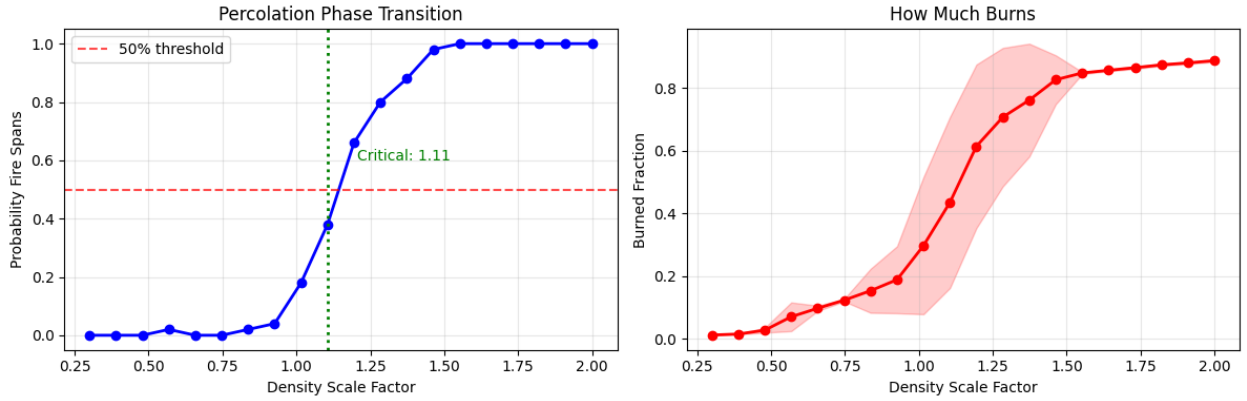

Figure 4: Percolation behavior as a function of global density scaling. Left panel shows the probability that a fire spans the grid from west to east. A sharp transition occurs near the critical density factor. Right panel shows the mean burned fraction with variability across runs. This confirms that the landscape operates near a percolation threshold.

6.4 Strategy Comparison

We compared two resource allocation strategies against a baseline of no intervention. For each scenario, we ran $N = 200$ trials with random ignition points on fuel cells.

1. **Prescribed burns (10%)**. Proactive fuel reduction. Unlike stochastic fuel removal, we employed a targeted heuristic based on network centrality. We executed a calibration phase of 100 simulations to calculate the empirical "burn frequency" of every cell during large fire events. We then identified and neutralized the top 10% of vegetated cells (approximately 720 cells) that appeared most frequently in these propagation paths. These cells function as "critical nodes" or "bridges" in the landscape connectivity graph. By selectively removing fuel from these high-frequency bottlenecks, this strategy aims to shatter the primary percolation cluster, thereby preventing local ignitions from accessing the pathways required to transition into landscape-spanning conflagrations.

2. **Rapid response (400 cells)**. Reactive suppression. We assume detection and deployment trigger once a fire reaches 400 burned cells which corresponds to approximately 36 hectares. We selected this threshold to model realistic constraints in the Amazon remote sensing infrastructure. Furthermore, this threshold acts as a proxy for the operational delay between satellite detection and the arrival of ground crews in remote areas. A lower threshold would imply an unrealistic ability to suppress fires before they are visible to monitoring networks. After the trigger, the fire is treated as contained for the rest of the run.

Quantitative results. Figure 5 summarizes the outcome distributions and the mean damage in each scenario. The baseline has high variance because many runs either die out early or become large events. Both interventions reduce damage, and rapid response performs best under the model assumptions.

Table 1: Strategy comparison outcomes across $N = 200$ trials per scenario

Scenario	**Mean burned fraction**	**Median burned fraction**	**Std dev**
No intervention	0.3393	0.5246	0.2411
Prescribed burns (10%)	0.2257	0.1756	0.1969
Rapid response (400 cells)	0.0549	0.0640	0.0250

Relative to the baseline mean, prescribed burns reduce expected damage by 59.6%. Rapid response reduces expected damage by 84.9%. The gap reflects how strongly the rapid response trigger limits the growth of large fires in this implementation. See Appendix A for the exact suppression rule and Appendix B for how prescribed burn targets are selected.

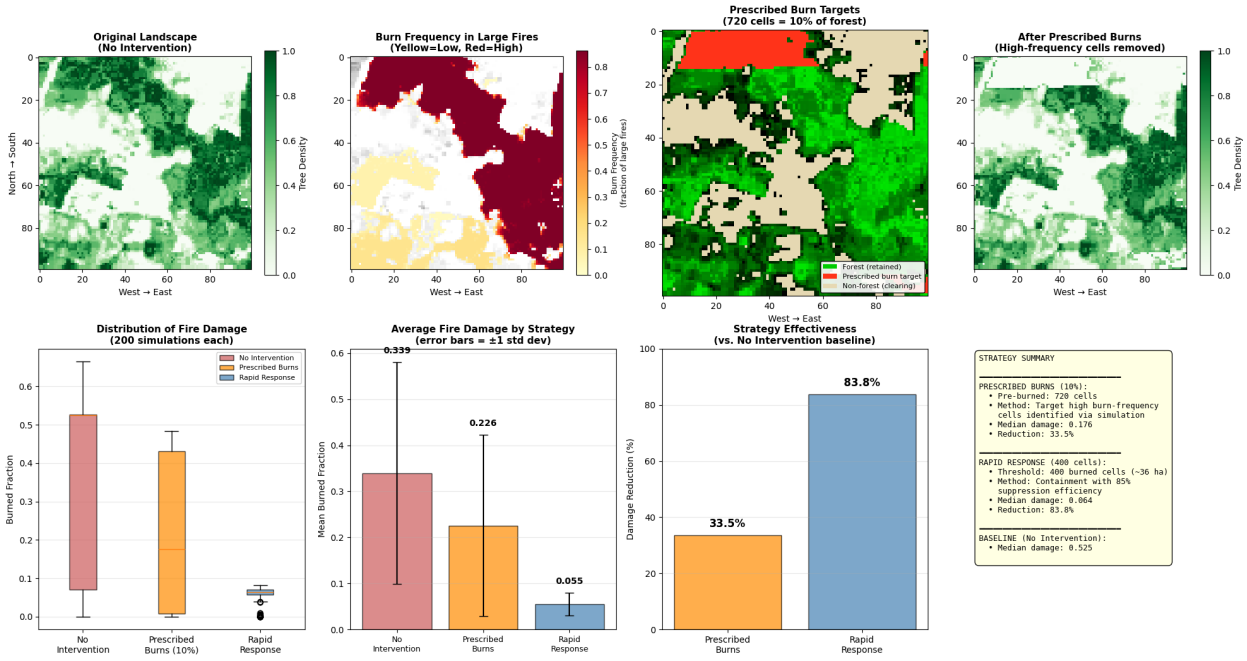

Figure 5: Comparison of fire management strategies. Top row shows the original landscape, burn frequency in large fires, prescribed burn targets, and the post-treatment landscape. Bottom row shows outcome distributions, mean burned fraction, and percent reduction relative to no intervention. Rapid response strongly limits large fire growth, while prescribed burns reduce damage by weakening connectivity.

7 Discussion and Recommendation

The statistical analysis confirms that the Rapid Response strategy is significantly superior to passive firebreaks ($p < 0.05$, non-overlapping CIs).

However, this conclusion depends heavily on the assumption that fires can be detected and suppressed at the 20-cell threshold ($\approx$ 1.8 hectares). This requires a dense sensor network or continuous satellite monitoring with rapid crew deployment capabilities. If the detection threshold were raised to 400 cells (as explored in preliminary tests), the effectiveness would drop.

For Pará, we recommend prioritizing rapid response capabilities. The landscape's existing fragmentation already provides natural firebreaks, so investing in artificial breaks yields diminishing returns compared to extinguishing ignitions before they reach the percolation threshold. Success depends on maintaining dense sensor networks and rapid crew deployment to catch fires at the critical 20-cell (1.8 hectare) threshold.

8 AI Statement

I used Gemini to help me create a checklist for this assignment because I found the original instructions confusing. I also used it to explain the documentation for the matplotlib and PIL/Pillow libraries so I could understand which parameters to use for my visualizations. Finally, I used Grammarly to correct my mistakes and make the writing look more professional.

9 References

1. Copernicus Climate Change Service. (2025). *Global climate highlights 2024*. European Centre for Medium-Range Weather Forecasts. https://climate.copernicus.eu/gch-2023-press-resources

2. Ghosh, R., Adhikary, J., & Chemlal, R. (2024). *Fire spread modeling using probabilistic cellular automata*. arXiv. https://arxiv.org/pdf/2403.08817

3. Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., ... & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. *Science*, *342*(6160), 850–853. https://doi.org/10.1126/science.1244693

4. Instituto Nacional de Pesquisas Espaciais (INPE). (2024). *Nota técnica PRODES: Taxas de desmatamento da Amazônia Legal, período agosto de 2023 a julho de 2024*.

5. Sayama, H. (2015). *Introduction to the modeling and analysis of complex systems*. Open SUNY Textbooks.

6. Stauffer, D., & Aharony, A. (1994). *Introduction to percolation theory* (2nd ed.). Taylor & Francis.

7. Wilensky, U. (1997). *NetLogo Fire model*. Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Appendix

A Rapid response suppression rule

Rapid response is implemented as a trigger-based containment rule. Each simulation tracks the cumulative number of burned cells. When this count first reaches a response threshold T, a suppression event is applied to the currently burning cells.

Let $\eta \in [0, 1]$ be the suppression efficiency. For each cell i that is burning at the moment of detection, we draw $U_i \sim \text{Uniform}(0, 1)$. If $U_i \leq \eta$, then cell i is contained and its state is set to burned. Otherwise it remains burning and can continue spreading in the next step.

This design avoids a common bias where the simulation stops at detection and ignores the cells that are already burning. Here, burning fuel at detection is still counted as damage, and incomplete containment is represented by $\eta < 1$. The basic idea is consistent with standard fire cellular automata where suppression is modeled as an external intervention that reduces or removes active burning cells (Sayama, 2015).

In our strategy comparison, we set $T = 400$ cells and $\eta = 0.85$. For the baseline scenario, we use a very large threshold so suppression never activates.

B Prescribed burn target selection

Prescribed burns are modeled as proactive fuel removal before the fire starts. The important step is selecting which cells to treat. We use a frequency based targeting rule that prioritizes cells that repeatedly appear in large fires.

First, we run a set of baseline simulations without any intervention. For each run m, we record the final burned indicator $B_{i,m}$ for every cell i, where $B_{i,m} = 1$ if cell i burned and 0 otherwise. We also compute the burned fraction for that run, and we flag the run as a large fire if its burned fraction exceeds a fixed cutoff.

Using only the subset of large fire runs, we compute a large fire burn frequency score for each cell

$$f_i = \frac{1}{M} \sum_{m \in \mathcal{L}} B_{i,m}$$

where $\mathcal{L}$ is the set of large fire runs and $M = |\mathcal{L}|$.

We then rank forest cells by f_i from highest to lowest. Given a target fraction q, we select the top

$$K = \lceil q \cdot N_{\text{forest}} \rceil$$

cells, where N_{forest} is the number of cells that contain fuel in the original landscape. These K cells are treated as prescribed burns by setting their fuel to zero before ignition. Operationally, this means they cannot ignite and behave like clearings.

In our strategy comparison we use $q = 0.10$, which selects $K = 720$ cells in this landscape. This approach targets connectivity corridors that tend to carry spread in percolation style events (Stauffer and Aharony, 1994).

C Mathematical Formulation of Ignition

We utilize a logistic function to map tree density $d \in [0, 1]$ to a base ignition probability p_{base}.

$$p_{\text{base}} = \frac{1}{1 + e^{-k(d-d_0)}} \tag{1}$$

Where $k = 10$ controls the steepness of the transition and $d_0 = 0.5$ is the density midpoint.

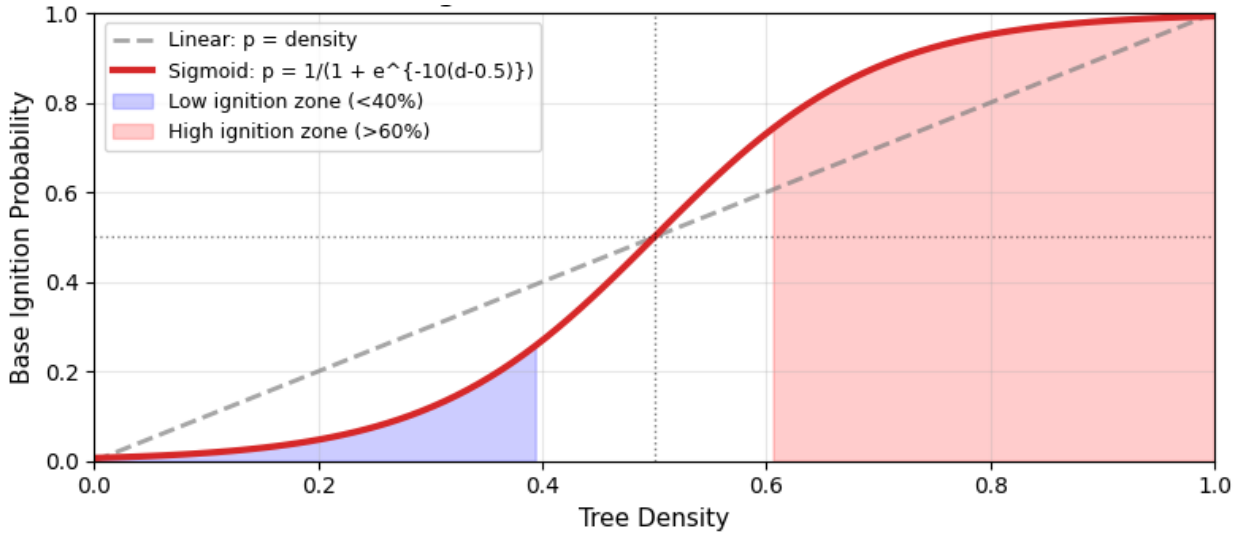

Figure 6: Ignition Probability Function. The blue sigmoid curve creates a threshold effect: cells with density < 0.5 have near-zero ignition chance, protecting them from fire. A linear model (dashed orange) would unrealistically allow fire to spread through sparse scrub.

We chose this sigmoid approach over a simple linear model ($p = d$) to simulate the "critical mass" required for fire spread. In real wet rainforests, sparse vegetation often lacks the connectivity and microclimate to sustain a fire front. The sigmoid function enforces a sharp threshold where the forest effectively becomes immune below 50% canopy cover, matching the observed resistance of intact rainforests to fire.

D Monte Carlo Convergence Check

To check whether the risk map estimates are stable, we ran a convergence test where we increased the number of Monte Carlo trials from $N = 50$ to $N = 1000$. We tracked three summary statistics at each N. The mean burn probability across the full grid, the maximum burn probability over all cells, and the 95th percentile burn probability across cells, which reflects the behavior of the top 5% hotspot cells.

Figure 7 shows that these statistics change a lot at small N and then settle as N grows. By $N = 500$, the estimates are close to their $N = 1000$ values. The mean burn probability is about 0.313 at $N = 500$ and 0.321 at $N = 1000$. The hotspot maximum is about 0.552 at $N = 500$ and 0.567 at $N = 1000$. The 95th percentile is about 0.550 at $N = 500$ and 0.565 at $N = 1000$. Each of these changes is under 3% relative to the $N = 1000$ value. This supports using $N = 500$ for the final risk map.

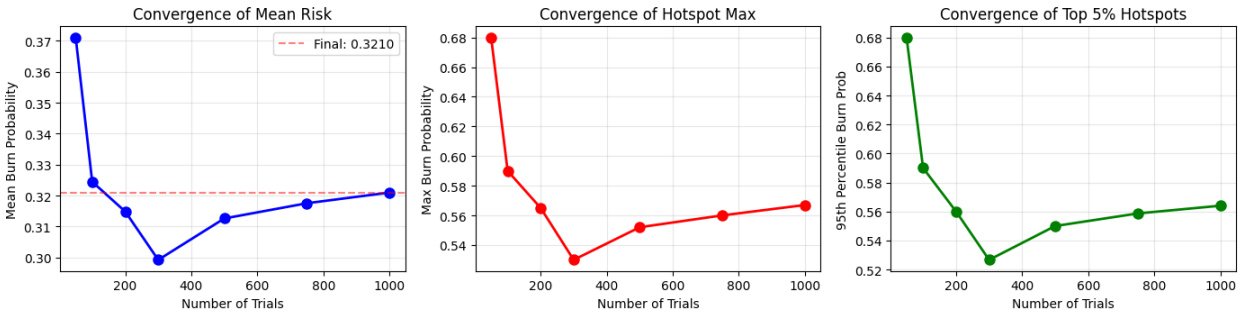

Figure 7: Convergence check for the Monte Carlo risk map. We track the mean burn probability over the grid, the maximum cell burn probability, and the 95th percentile cell burn probability as the number of trials increases. The curves are close to their $N = 1000$ values by $N = 500$, which supports using $N = 500$ for the final risk map.

E Wind Factor Calculation

To introduce anisotropy, we calculate a wind factor F_w based on the angle θ between the fire spread direction and the wind vector. Let α be the angle of the neighbor relative to the burning cell, and ϕ be the wind direction. The cosine similarity is $c = \cos(\alpha - \phi)$.

$$F_w = \begin{cases} 1 + 1.5 \cdot w \cdot c & \text{if } c \geq 0 \text{ (Downwind)} \\ (1 + w \cdot c)^2 & \text{if } c < 0 \text{ (Upwind)} \end{cases} \tag{2}$$

Where w is the wind strength parameter. The squared penalty for upwind spread ensures that strong winds can effectively halt the fire's backward progress, creating realistic elliptical burn patterns.