Machine Learning with Applications 22 (2025) 100779

Contents lists available at ScienceDirect

Machine Learning with Applications

journal homepage: www.elsevier.com/locate/mlwa

Identifying critical fire spread to the wildland–urban interface using cellular automata and reinforcement learning

Javier González-Villa [a,*], David Lázaro [a], Arturo Cuesta [b], Adriana Balboa [b], Daniel Alvear [b], Mariano Lázaro [b]

[a] Department of Applied Mathematics and Computer Science, University of Cantabria, Spain
[b] Transportation and Project and Process Technology Department, University of Cantabria, Spain

ARTICLE INFO

Dataset link: https://zenodo.org/records/14991404

Keywords:
Wildfire
Cellular automata
Stochastic simulation
Reinforcement learning
Wild-land urban interface
Rothermel surface fire spread model

ABSTRACT

Wildfires threatening the wildland–urban interface present significant risks to community safety, especially under conditions of inadequate vegetation management and adverse weather. Accurately identifying scenarios in which fire reaches this interface is crucial for timely evacuation planning and risk mitigation. This study presents a computational method using cellular automata and stochastic simulations to model wildfire spread. Stochastic scenarios generated through the cellular automata are employed to train a reinforcement learning model, which leverages computer vision techniques to interpret multiple layers representing diverse environmental factors. This enables the reinforcement learning agent to identify and prioritise critical fire trajectories that could impact the wildland–urban interface. The framework adapts the Rothermel surface fire spread model within a cellular automata structure, providing a simplified yet effective simulation of fire propagation under variable conditions. The proposed approach was validated using synthetic and real-world case studies, demonstrating its potential for integration with geographic information systems. Results suggest this approach enhances the identification of critical fire spread scenarios and improves computational efficiency for real-time applications. By enabling real-time recognition of high-risk events, our framework supports more informed evacuation strategies and fire management decisions around the wildland–urban interface.

1. Introduction

Wildfires represent a significant challenge to ecosystems, biodiversity, economy, and human security on a global scale. Their frequency and severity are increasing owing to climate change, inadequate forest management and other human factors (Thompson & Calkin, 2011) which require advanced prediction and response tools. Understanding and predicting wildfire behaviour is essential for effective environmental management and the development and applications of public safety measures. Wildfire spread models, which are sophisticated mathematical and computational tools, play a key role in simulating fire dynamics and could be valuable for informed decision making and the development of mitigation and response strategies. These models can be classified into empirical models, which are based on relatively simple analytical equations experimentally defined, such as the fire spread velocity defined in the Rothermel model (Rothermel, 1972) (e.g. FARSITE: Fire Area Simulator—Model Development and Evaluation, USDA) or physical models that use physical, chemical, and heat transfer principles to represent ignition and fire behaviour, calculate spread rates, fire intensity, and energy production, among other parameters (e.g. FIRETEC (Linn & Harlow, 1997) or WFDS (Mell, 2010)). Nevertheless, there is still an issue in the consideration of the inhomogeneity of the input parameters of the models, such as vegetation and fire-atmosphere, which can significantly affect fire propagation, even transitioning to chaotic propagation (Danta, Egorova, & Pagnini, 2024). In furtherance of their intended core functions, these models contribute significantly to the refinement of existing global evacuation models such as EMS (González-Villa, Cuesta, Alvear, & Balboa, 2022) and WUI-NITY (Wahlqvist et al., 2021) and similar models that provide comprehensive evacuation plans and predictions of required evacuation times for natural disasters (Gwynne et al., 2023). Providing a prediction of the arrival time of the flame front at the wild-land urban interface allows a comparison between the Available Safe Evacuation Time (ASET) and the time required for evacuation, as estimated by these models. This improves the accuracy and reliability of safety measures

* Corresponding author.
E-mail addresses: javier.gonzalezvilla@unican.es (J. González-Villa), david.lazaro@unican.es (D. Lázaro), arturo.cuesta@unican.es (A. Cuesta), adriana.balboa@unican.es (A. Balboa), daniel.alvear@unican.es (D. Alvear), mariano.lazaro@unican.es (M. Lázaro).

https://doi.org/10.1016/j.mlwa.2025.100779
Received 21 July 2025; Received in revised form 20 October 2025; Accepted 30 October 2025
Available online 3 November 2025

and ensures timely and efficient evacuation strategies during forest fires. Nevertheless, to be applied in this way, wildfire spread models should comply with two main requirements: real-time simulation and safety results.

In this regard, stochastic simulation models for both evacuation and fire propagation prediction purposes have become essential tools in wildfire risk management, allowing for the exploration of a range of potential fire scenarios and consequences under varying scenario conditions (Beneduci & Mascali, 2024; Boychuk, Braun, Kulperger, Krougly, & Stanford, 2009; Buch, Williams, Juang, Hansen, & Gentine, 2023; González-Villa et al., 2022; Hajian, Melachrinoudis, & Kubat, 2016; Nguyen, Schlesinger, Han, Gür, & Carlson, 2019; Waeselynck & Saah, 2024). These models rely on probabilistic representations of key factors such as weather patterns, fuel characteristics and ignition locations in the event of a wildfire, or population densities and population characteristics in the evacuation scenario to generate sets of simulations. By analysing the statistical properties of these ensembles, it is possible to quantify the likelihood of different fire and evacuation outcomes and assess the potential impacts on communities and infrastructure.

Recent developments have highlighted the potential of machine learning algorithms to predict wildfire behaviour in real time and to optimise input parameters for physics or empirically based wildfire spread models to improve fire spread predictions (Jain et al., 2020a), marking a promising path for the integration of more advanced artificial intelligence techniques. Namely, the application of Deep Reinforcement Learning and Cellular Automata, has gained traction in wildfire spread modelling. These approaches have shown promise in improving prediction accuracy and computational efficiency. For instance, through the use of spatial reinforcement learning to build forest wildfire dynamics models from satellite images (Ganapathi Subramanian & Crowley, 2018), the exploration of evolutionary cellular automata for modelling wildfires (Green, DeLuca, & Kaiser, 2020), or combined machine learning-based cellular automata with Geographical Information Systems (GIS) for forest fire spread modelling for diverse purposes (Murray et al., 2024; Sun et al., 2024).

In addition, innovative research using remote sensing technologies has greatly improved our ability to monitor and analyse the spatial dynamics of wildfires in real time, providing crucial data for model validation and improvement, which produces great synergies with the above-mentioned learning technologies (Leblon, Bourgeau-Chavez, & San-Miguel-Ayanz, 2012). Despite these advances, current models still have difficulties predicting the probabilistic direction of fire spread and integrating real-time response mechanisms, indicating a deficit in contemporary approaches to wildfire management in terms of purely population-based evacuation (Beyki, Santiago, Laím, & Craveiro, 2023). Deterministic models are inadequate for encapsulating the unpredictable nature of wildfires, which argues for hybrid methodologies that amalgamate the deterministic and probabilistic components (Singh et al., 2024).

In the Probabilistic Risk Assessment of other critical fire propagation scenarios, such as nuclear power plants, two concepts have been used in recent decades. These are the Maximum Expected Fire Scenarios (MEFS) and the Limiting Fire Scenarios (LFS) (Association et al., 2006). The MEFS was used to determine the expected fire propagation based on realistic input parameters and conditions to define the model. LFS results in unfavourable consequences with respect to performance criteria. That is, the LFS is based on a maximum possible, though unlikely, value for one input variable or an unlikely combination of input variables that allow the damage limits to be surpassed. These concepts are widely used in nuclear power land can be applied to wildfire management.

This approach aims to address the gap in the application of wildfire scenarios simulated by stochastic models as training datasets for reinforcement learning models to identify critical scenarios near the Wildland–Urban Interface (WUI) in real time. Using this framework, we aim to obtain predictions for ASET by incorporating MEFS and LFS. To achieve this, we adopt a stochastic simulation method to predict wildfire spread probability distributions by accurately emulating MEFS using cellular automata. These simulations are further enhanced with detailed spatial information provided by GIS integration, as performed in Velasquez, Munoz-Arcentales, Bohnert, and Salvachúa (2019).

Furthermore, using these cellular automata as environments, we integrated deep reinforcement learning techniques to probe LFS. This strategy aligns with recent European initiatives advocating comprehensive wildfire management tactics and is poised to make significant contributions to this field (Lier et al., 2022). To implement and evaluate the foundational spread models, we employed a technique similar to established wildfire spread models, such as FARSITE (Finney, 1998) or ELMFIRE (Lautenberger, 2013), running a simplified scenario based on the Rothermel surface fire spread model, adjusted for wind and slope effects. This model was digitally represented by cellular automata, allowing seamless integration with GIS data and digital twins, thus linking cartographic, climatological, and fuel model data across dynamic spatial scales. Rigorous stochastic simulations and model training iterations are performed throughout the simulation phase to identify the most probable and critical fire scenarios.

This integration of immediate and long-term perspectives provides fire management professionals with a powerful framework for developing more effective strategies to combat the increasing threat of wildfires at the WUI. Furthermore, the ability of the model to simulate complex scenarios with multiple fire fronts and varied environmental conditions enhances its utility for evacuation planning. Identifying potential bottlenecks and high-risk areas in evacuation routes allows emergency managers to optimise evacuation strategies and allocate resources more effectively. This comprehensive approach not only improves the accuracy of evacuation time estimates but also contributes to the development of more resilient communities in fire-prone regions.

2. Methods

Given the objective of surface fire modelling, our development process adhered to a specific methodology, similar to that of widely validated classical comprehensive wildfire spread models, such as FARSITE (Finney, 1998), PhyFire (Asensio et al., 2021), FIREMAP (Vasconcelos & Guertin, 1992) or ELMFIRE (Lautenberger, 2013) but with some reductions. These reductions include the influence of phenomena such as canopy spread, spotting, or smouldering, which have been extensively studied in other research (Albini, 1979; Albini, Alexander, & Cruz, 2012; Liu, Lei, Gao, Chen, & Xie, 2021; Purnomo, Christensen, Fernandez-Anez, & Rein, 2024) but are of no relevance to our model given the specific objective and WUI scenarios considered. Although canopy spread, spotting, and smouldering influence some wildland fires, their role in many WUI surface fire events is often secondary compared to surface fire dynamics (Morvan, 2011). This simplification is further justified in the context of this model's application given that in Spain, the large majority of WUI vegetation consists predominantly of scrubland and agricultural land, rather than extensive forest canopy. According to the Ministerio para la Transición Ecológica *y* el Reto Demográfico, approximately 60%–65% of combustible vegetation near WUI areas is low-height scrub and brush, while cultivated land and pastures account for approximately 25%–30% (MITECO, 2023). Dense forest cover, characterised by a continuous canopy that favours crown fires, represents only a small fraction of these areas. This has also been identified by studies using LIDAR and GIS data confirming these combustible material arrangements, where it can be seen that forest mass represents less than 15% of land cover in WUI areas (Olmo, Ruiz, Planelles, & Hernando, 2013). Excluding these processes enables computational efficiency necessary for real-time applications. However, it should be noted that in particular WUI settings — especially those with dense, continuous canopy or extreme wind — the omission may lead to an underestimation of fire intensity or rate of spread (Cruz, Alexander, & Wakimoto, 2004) . Specifically, we used the Rothermel surface fire

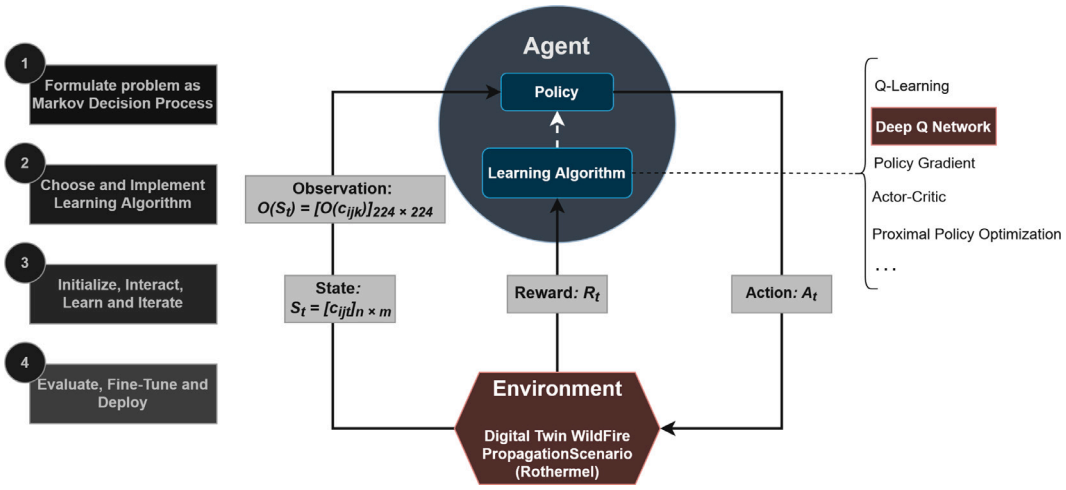

Fig. 1. Reinforcement learning approach adopted.

spread model (Andrews et al., 2018) with some adjustments proposed by Andrews, Cruz, and Rothermel (2013) and Anderson (1983) in terms of fitting the wind and slope influence parameters.

Cellular automata provide a computational means to represent fire scenarios with precise cell size and adaptative resolution. This approach allows alignment with digital twins, GIS data, and satellite-derived data resolutions, which facilitate seamless integration. This model also serves as a vital bridge between traditional propagation models, computational techniques, and a deep reinforcement learning approach (Sutton & Barto, 2018), acting as a computational interface for their interaction. This environment allow us to efficiently conducted probabilistic mapping of fire spread under diverse meteorological scenarios following a stochastic approach.

The capability provided by cellular automata to represent scenarios also acts as an environment to train and validate a smart agent exploring the applicability and efficiency of Deep Q-Networks (DQNs) through a reinforcement learning approach (Fig. 1) to identify critical scenarios where climatic and scenario boundary factors converge. This strategic approach enables prompt intervention and evacuation planning in response to potential wildfire disasters, allowing estimation of the Available Safe Egress Time (ASET) (Hurley et al., 2015) under MEFS and LFS.

2.1. Rothermel wildfire propagation model

Regarding the wildfire propagation model used, the Rothermel surface fire spread model (Rothermel, 1972) was adopted because it is a widely used quasi-empirical model for predicting the wildfire spread rate. The model is based on the heat balance model developed by Frandsen (1971) and the data obtained from wind tunnel experiments in artificial fuel beds with varying characteristics (Rothermel & Anderson, 1966). It also incorporates Australian wildfire data from grasslands (McArthur, 2000). The basic Rothermel model with minor modifications proposed by Albini (1976) has been used in the same form for decades. The key equation of the Rothermel model is as follows:

$$R = \frac{I_R \xi (1 + \phi_w + \phi_s)}{\rho_b \epsilon Q_{ig}} \tag{1}$$

The Rothermel equation establishes the ratio between the propagating heat flux, defined in the numerator, and heat sink, defined in the denominator. The basic propagating heat flux without considering the slope nor wind effects is defined as the reaction intensity I_R multiplied by the propagation flux ratio ξ. Subsequently, additional convective and radiative heat due to the wind and slope can be considered with the dimensionless variables ϕ_w and ϕ_s. These variables, although complex to calculate given their adjustments for the several fuel models, depend directly on the packing ratio, the surface area to volume ratio and on the slope and wind speed directly. For a more detailed breakdown of the formulas, given their complexity, please refer to Table 3 of Andrews et al. (2018) The heat sink defines the heat required to ignite the fuel, and it is defined by the heat of preignition Q_{ig}, the bulk density ρ_b and the effective heating number ϵ, which are well detailed in the original model (Andrews et al., 2018) and are related mainly with the different fuel models and factors like moisture content M_f, dead fuel moisture of extinction M_x, mineral content factor η_S, fuel depth δ, optimum reaction velocity r', net fuel loading w_n and low heat content h. Below is an expression of these factors related to the intensity of the reaction, detailing the effect on the calculations of the moisture content of the fuel.

$$I_R = r' w_n h \eta_M \eta_S \tag{2}$$

$$\eta_M = 1 - 2.59 \frac{M_f}{M_x} + 5.11 \left(\frac{M_f}{M_x} \right)^2 - 3.52 \left(\frac{M_f}{M_x} \right)^3 \tag{3}$$

This model also operates under the assumption of a linear flame front and computes the propagation speed of a head fire, considering the impact of wind and slope, as mentioned previously. this relies on the presence of fine dead fuel, which is a critical component of the fuel mix. This model is tailored for analysing fires in surface fuels close to the ground, excluding scenarios such as canopy fires in tall trees or slow-burning ground fires. Notably it can affect diverse landscape features, fuel properties, and weather conditions. It is adaptable to a range of fire spread mechanisms, even those of a complex nature.

In addition to the initial Rothermel model, some modifications proposed in Andrews et al. (2013) and Anderson (1983) were also considered where the midflame wind speed was obtained from the free wind speed and fuel bed depth following an unsheltered wind adjustment factor approach (Andrews, 2012). This approach is very convenient because only surface propagation is considered. Thereafter, the midflame wind speed is limited to a maximum value following the next expression $w_l = 96.8 \cdot I_R^{1/3}$ considering its relation to the reaction intensity. In addition, the effects of increasing the free wind speed with respect to altitude and terrain type were considered, following Hellmann's Exponential Law:

$$w_h = w_{10} \cdot (\frac{h}{10})^{\alpha} \tag{4}$$

where the wind speed at altitude h, w_h is calculated based on the wind speed measured at an altitude of 10 m w_{10}, considering the altitude h and terrain factor α that depends on the roughness and type of terrain (see Bañuelos-Ruedas, Angeles-Camacho, and Rios-Marcuello (2010)).

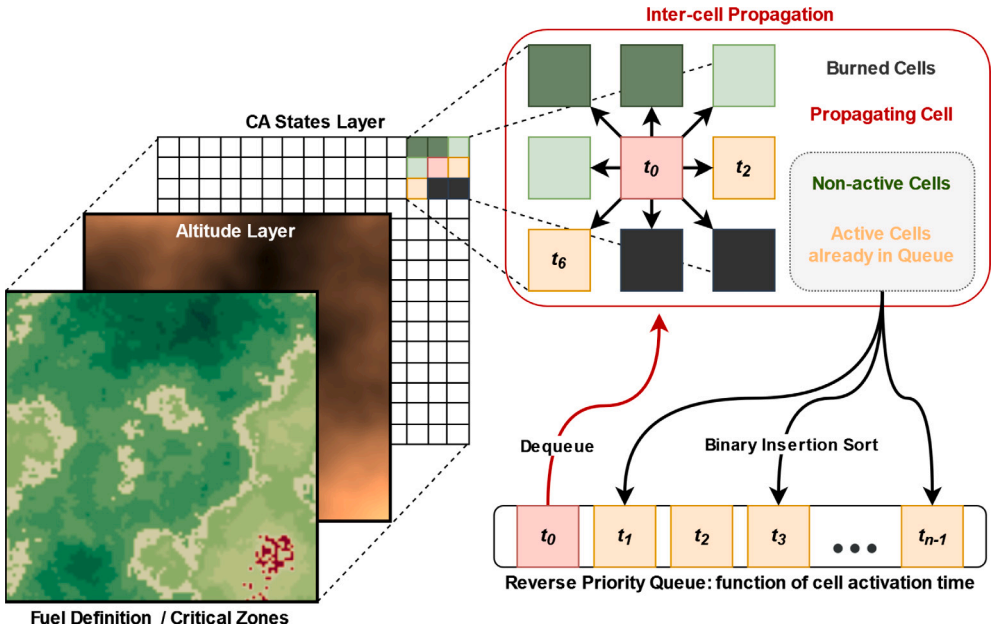

Fig. 2. Cellular automaton fire propagation methodology and information layers used.

The influence of wind direction on the propagation direction was also considered. For this purpose, the classical approach has been followed, in which the influence is represented elliptically (Anderson, 1983) with the ratio between the length and width of the ellipse depending on the front and back spread rates as factors that model the ellipse in terms of the definition function, location of the propagation focus within it, and elongation of the ellipse. Thus, this factor can reduce the influence of the wind on lateral propagation, allowing the precise determination of the influence of a given wind forming a specific angle of application with respect to the current direction of propagation of the flame front.

To implement this model computationally, each cell within the cellular automata is characterised by environmental data obtained from its geographic position, including fuel models and altitude. Wind data can be obtained from the nearest weather station and adjusted for altitude differences to estimate wind conditions in neighbouring cells using the described adjustment in Eq. (4). Once the characteristics of each cell have been established, fire spread begins by activating adjacent cells according to the Rothermel model, calculating slopes using altitude differences and calculating the speed of fire spread based on the parameters contained in the cell used to calculate Eq. (1).

2.2. Environment: Cellular automata

Once the propagation of our problem was studied, the next step was to find a nexus instrument between technologies. The selected instrument was cellular automata, a well-established mechanism already used in prior research (Boychuk et al., 2009; Freire & DaCamara, 2019; Ramírez, Lendechy, Castillo, Toscano, Medina, Vargas, Carballo, Romero, & Rosas, 2016). This modelling methodology is well suited for establishing a training environment with a specific spatial size $n \times m$ and T time steps for a deep reinforcement learning agent, providing the generation of stochastic environments that adhere to coherence parameters and physical principles governing the spread of surface wildfires, as well as enabling the integration of real data obtained from GIS with varying spatial and temporal resolutions, as substantiated by previous investigations (Green, Finney, Campbell, Weinstein, & Landrum, 1995).

The cellular automata model is meticulously designed to satisfy specific rules, allowing it to be represented as a Markov Decision Process (MDP) with the set of cells denoted as a tensor $C = [c_{ijt}]_{n \times m \times T}, c_{ijt} = (ft, h, M_f, s)$ where h is the altitude, s is the current status of the cell (burned, propagating, non-active or active) and ft is the set of attributes of the fuel model related to Eq. (1). This representation empowers the DQN-Agent model to make informed decisions and execute actions based on observations $O(c_{ijt}) = (ft, h, s)$ within the simulated environment and also enables the application of the stochastic approach to the identification of likely scenarios. The fire propagation within the cellular automaton is orchestrated such that each cell could spread to its eight adjacent cells (see Fig. 2), by calculating the activation times as a function of the fire propagation speed, as shown in Eq. (1), according to the Rothernmel model presented above and the characterisation of cells across terrain and biomes represented in terms of vegetation species and distribution, fuel density, altitude and moisture content, which are considered in the presented spread model (see Eq. (2) and (3)). As illustrated in Fig. 2, the cells can exist in four distinct states. The first is the propagating cell, which represents the active cell currently undergoing evaluation for its potential to spread. Second, there are active cells in the queue, comprising all the remaining active cells awaiting assessment by the model. The third state consists of the burned cells, that have already undergone combustion. Finally, the inactive cells, remained unaffected by the propagation process.

The activation time serves as the critical variable for determining the priority in a reverse-order queue, which governs the sequence of cell propagation. Each cell is evaluated for altitude, slope, wind speed and direction, moisture, and fuel properties. For this preliminary approach, while acknowledging the potential for expansion to encompass a broader range of fuel models (Andrews et al., 2018), we selected a subset of 12 fine dead fuel models from literature (Andrews et al., 2018) detailed in Table 1.

The Rothermel model with the above adjustments was used to simulate the spread of the fire surface, recording the activation and complete combustion times of each cell. These values are used later in the expected probabilistic calculation of wildfires. The characterisation parameters for each cell contain the intrinsic parameters of the fuels, delineated in Table 1 and linked to the Rothermel surface combustion model, and the external factors like the altitude for slope computations with neighbouring cells and wind factor relative to the 10-metre altitude level in accordance with the Hellmann Exponential Law. These parameters were stochastically generated considering coherence, mitigating abrupt alterations in the terrain, and ensuring vegetation uniformity. The Python Perlin Noise library was harnessed to generate procedural maps via gradient noise, smoothly interpolating across a pseudo-random matrix of values (Etherington, 2022), and

Table 1
Fuel models used to characterise each cell with particular vegetation for the Rothernmel wildfire surface propagation model.

Fuel model [code]	Surface area - volume ratio (σ) [m^{-1}]	Oven-dry fuel load (w_0) [kg/m^2]	Fuel depth (δ) [m]
No Fuel [0]	–	–	–
Grass (short) [1]	11 483	0.1660	0.3048
Grass (tall) [2]	4921	0.6738	0.7620
Brush (not chaparral) [3]	6562	0.2246	0.6096
Chaparral [4]	6562	1.1230	1.8288
Timber (grass and understory) [5]	9843	0.4492	0.4572
Timber (litter) [6]	6562	0.3369	0.0609
Timber (litter and understory) [7]	6562	0.6738	0.3048
Hardwood (litter) [8]	8202	0.6542	0.0609
Logging slash (light) [9]	4921	0.3369	0.3048
Logging slash (medium) [10]	4921	0.8984	0.7010
Logging slash (heavy) [11]	4921	1.5720	0.9144

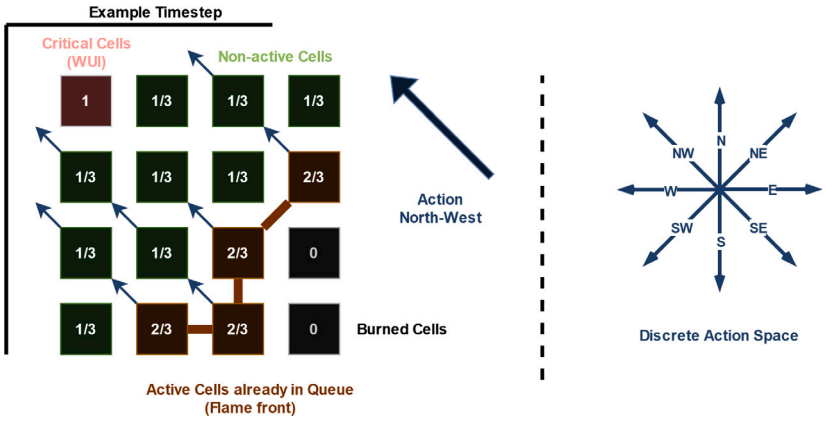

Fig. 3. Example visualisation of the action space and the experience buffer by means of a timestep representation of one of the state channels of the cellular automaton.

serving as the foundation for producing coherent patterns as illustrated in Fig. 2. The adeptness of the cellular automaton model in generating coherent maps and managing specific wind conditions facilitates its direct engagement with decision-making agents. We will see below that these capabilities are very adaptive for various purposes and interact very well with MEFS and LFS computations.

2.3. MEFS: Stochastic approach

MEFS plays a key role in managing mass evacuations near the WUI. Our approach employs a sophisticated stochastic model based on a cellular automata framework to characterise wildfire behaviour across diverse landscapes. This model incorporates both fixed variables (terrain slope, vegetation conditions, and ignition points) and dynamic variables (wind direction and velocity), which are adjusted during the simulation time steps based on historical data or random distributions. Our stochastic approach significantly expands the capabilities of MEFS beyond providing a single estimation of the most plausible scenario. While real-time predictive capability is crucial for immediate response, our model versatility allows for a comprehensive risk analysis through the introduction of multiple ignition points and diverse climatic combinations. This enhanced MEFS approach enables the following:

- The generation of probabilistic risk maps that identify areas that are most susceptible to fire propagation offers a nuanced understanding of landscape vulnerability.
- Sensitivity analysis of fire behaviour in response to variations in initial conditions and environmental parameters provides insights into the robustness of predictions.
- Simulation of complex scenarios involving multiple simultaneous fire fronts, mirroring high-risk situations encountered in real-world wildfire events.
- Provision of valuable data for long-term strategic planning in forest management and risk mitigation, allowing for proactive rather than reactive approaches.

To produce these estimations, multiple fire simulations denoted as N must be conducted on a consistent landscape, tracking the fire propagation times and affected cells. These simulations were performed under the same scenario, keeping the static variables constant, and varying the wind direction and speed at each instant of the simulation. This approach offers two possible methodologies, depending on the objective pursued: (1) variation according to historical data or weather forecasts that are used to predict the MEFS in specific scenarios, or (2) complete wind variation in the action space (Fig. 3) in all possible directions considering multiple ignition points in this scenario, which is used for the global risk assessment of an entire area without specifying a specific scenario.

By aggregating these spread times results $\tau = [\tau_{ijk}]_{n\times m\times N}$, we calculated the average and minimum values for different purposes. These values are used to develop a probability map $P = [p_{ij}]_{n\times m}$ illustrating the likelihood of cell activation $p_{ij} = \frac{1}{N}\sum_{k=1}^{N} I(\tau_{ijk} < \mu)$ considering a threshold μ which would correspond to the required safe evacuation time from an area near the WUI. Also, the $LFS = [lfs_{ij}]_{n\times m}$ provided by the stochastic model (MST) is calculated via $lfs_{ij} = \min_{k\in\{1,\dots,N\}}\{\tau_{ijk}\}$ to identify the most critical scenario. This stochastic method captures wildfire uncertainties and provides strong fire spread risk estimates near the WUI. The resulting probability map and estimated ASET allows for targeted, scenario-specific simulations based on real or predicted weather patterns, as well as broader, risk-assessment simulations that consider all potential wind conditions.

This versatility enhances the applicability of the model to various fire management and planning contexts.

Moreover, it also plays a pivotal role in training more advanced LFS models, which aim to surpass the stochastic model by simulating the most unfavourable conditions leading to the most dangerous fire scenarios in terms of ASET for WUI zones. By leveraging this expanded MEFS concept, our stochastic model serves as a robust decision-support tool, bridging the gap between immediate fire response and long-term preventive planning. It offers a holistic approach to wildfire risk assessment by combining the urgency of real-time prediction with the foresight of comprehensive risk analysis.

2.4. LFS: DQN-agent

The Deep Q-Networks (DQNs) approach is employed to identify Limiting Fire Scenarios (LFS) in wildfire spread that are more critical than those generated by the stochastic model, while maintaining computational feasibility. The DQN training relies on the Bellman equation, which governs the temporal difference updates of Q-values representing expected rewards for actions within given states.

$$Q(s,a) = r + \gamma \cdot \max_{a'} Q(s',a') \tag{5}$$

Where $Q(s,a)$ denotes the Q-value for state s and action a, r is the immediate reward, γ is the discount factor gamma determining the influence of future rewards. $\max_{a'} Q(s',a')$ is the maximum Q-value for the next state s' and all possible actions a'. For each training step, the neural network predicts the Q-values for a given state. An action is then selected based on these Q-values using an epsilon-greedy policy to balance exploration and exploitation. After executing the action and observing the reward and new state, the target Q-value is calculated using the Bellman equation. The difference between this target value and the Q-value predicted by the network is used as an error signal to update the neural network weights. The loss function, optimised using the Adaptive Moment Estimation (Adam) method used to train the network, is the mean squared error between the predicted and target Q-values:

$$L = (r + \gamma \cdot \max_{a'} Q(s',a') - Q(s,a))^2 \tag{6}$$

This approach is particularly suitable for problems with large state spaces and complex action sequences, where traditional value matrices become impractical because of their combinatorial complexity. The DQN learns to process high-dimensional inputs, such as raw images, and makes decisions that maximise the expected future rewards based on the observed environment. The core of our methodology is a cellular automaton with Moore's neighbours, which serves as the training environment. This environment is governed by an MDP that closely models wildfire propagation dynamics based on the Rothermel model. The state space is well defined, with actions in each timestep A_t limited to selecting wind directions (north, south, east, west, northeast, northwest, southeast, and southwest) at each time step (see Fig. 3). The observation space consisted of three layers representing the cellular automaton: the altitude layer, containing continuous values from 0 to 1; the vegetation layer, incorporating discrete values from 0 to 1; and the cell state layer, featuring four distinct states: burned (0), inactive (0.333), active (0.666), and critical (1).

These layers are structured similarly to normalised colour channels in a composite image, thus allowing us to utilise the ResNet18 architecture, pre-trained on ImageNet and COCO datasets, partially using transfer learning since the layers and their weights are not frozen but allowed to fine-tune during the training process. The input images, with dimensions equal to the cell resolution of the automaton, were normalised and resized to 224 × 224 pixels using bilinear interpolation to ensure compatibility with the model requirements (see Fig. 5). The final network layer outputs eight Q-values, each linked to an action in the action space. Through experimentation, we determined that retraining all layers of the ResNet18 model yielded optimal performance, rather than selectively freezing some layers.

Table 2
Summary of model and training hyperparameters.

Parameter	Description
Optimiser	Adam
Loss function	Mean Squared Error (MSE)
Discount factor	$\gamma = 0.99$
Epsilon strategy	Epsilon-greedy decay $\epsilon_t = \max(0.01, \epsilon_0 \cdot \lambda^t),\ \epsilon_0 = 1,\ \lambda = 0.995$
Learning rate	$\alpha = 1e^{-3}$
Batch size	2^6
Replay memory	2^{12}
Reward	+1 (goal), −1 (fail)

The training process (see Fig. 4) involved balancing exploration and exploitation to identify the most effective decision sequence. The primary objective was to propagate the fire front towards the WUI, represented by the critical cells, in the shortest possible time. The reward R_t structure is designed to maintain a constant and controlled learning factor by providing a reward for successfully reaching the goal (reaching the WUI), along with a penalty for exceeding the time limit (minimum achieved by stochastic model). It is important to note that the optimal Q-values for each state, which are used for training the DQN model, are based on simulated wind data during the training phase and a combination of simulated and historical data for testing and validation stages. To enhance the learning experience, the agent is exposed to a variety of stochastic scenarios, including diverse, synthetically generated fire situations where the time limit corresponds to the minimum achieved by the purely stochastic model in multiple iterations. This exposure allows the agent to assimilate and comprehend the multifaceted impacts of various environmental variables, including wind direction, topographical gradients, and heterogeneous composition of vegetation types. The agent's performance was measured through two metrics: the number of consecutive times the agent was able to outperform the stochastic model's time and the relative time improvement against the same target time. The training was conducted on a GPU computing partition of the Altamira Supercomputer (RES), utilising CUDA technology to facilitate parallel processing of the complex calculations inherent in this advanced deep learning application. The DQN agent was trained using synthetic datasets comprising a total of 10,000 randomly generated scenarios (15 × 15 cells with 300 m spatial resolution) with varying fuel patterns, topographies and wind boundary conditions to ensure maximal diversity. Each training episode was initialised with a central ignition point and concluded either upon the fire reaching the WUI or when a maximum timestep limit, established as the minimum time achieved by the stochastic model, was met. The reward structure assigned a positive value when the agent successfully reached the critical WUI cells within this threshold, and a negative value otherwise. The configuration of hyperparameters — including discount factor, epsilon-greedy decay schedule, learning rate, batch size, replay memory size, and network update policy — can be found in Table 2. This comprehensive training scheme effectively combines reinforcement learning and stochastic modelling techniques, producing an agent capable of identifying critical wildfire scenarios and estimating the associated ASET.

3. Materials

Model verification scenarios. To verify the cellular automata implementation incorporating the Rothermel model for wildfire spread prediction, a set of scenarios was defined and systematically compared against analytical solutions. This verification process ensures the accurate representation and application of equations and numerical solvers within the cellular automata framework. The design of these scenarios also considered adjustments for wind influence and slope effects,

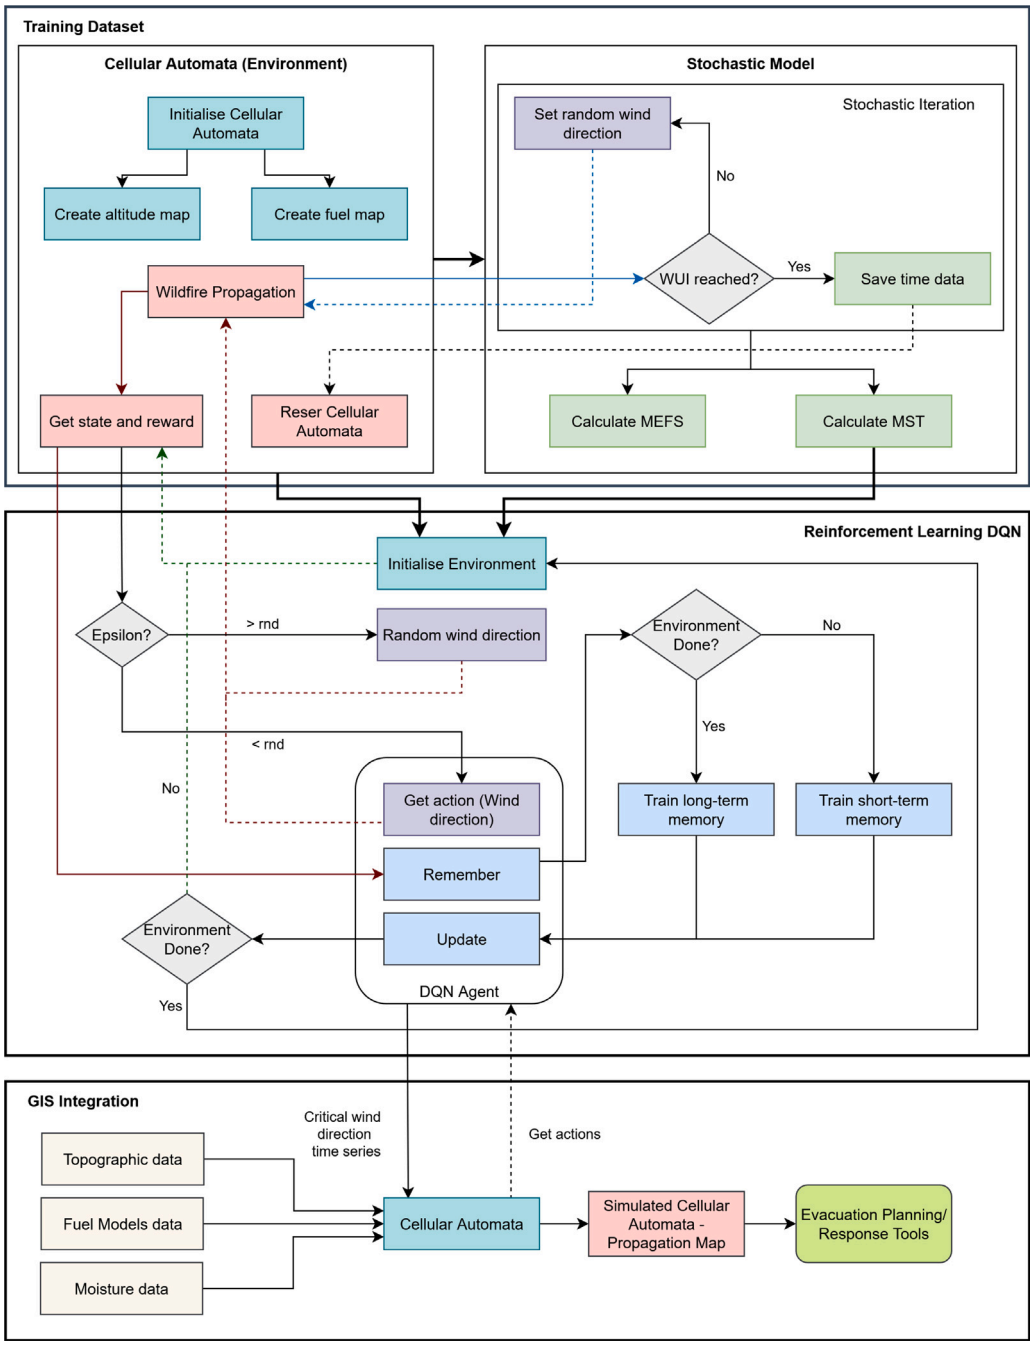

Fig. 4. Workflow between framework elements.

particularly their linkage to variables such as altitude and propagation direction. Furthermore, all verification scenarios were employed to conduct a preliminary sensitivity analysis and computational cost evaluation of the model as a function of spatial resolution and number of cells, providing foundational insights into the trade-offs between accuracy and efficiency in large-scale wildfire simulations. The verification scenarios (see Table 3) considered a surface fire spread area of 500 m × 500 m, which was defined in cellular automata with 500 × 500 square cells with a spatial resolution of 1 m and a spread time limit of 1 h. As shown in Table 3, the verification considered variations in the fuel model, slope, wind, and moisture content (fraction) with respect to the extinction moisture content.

Synthetic Training/Testing Scenarios. The second set of scenarios consists of synthetic scenarios of smaller size in terms of number of cells (15 × 15) randomly generated and designed for the purpose of training the DQN-Agent quickly and efficiently. These random scenarios are generated in a coherent way, in terms of terrain continuity and coherence of the vegetation distribution, and the zones that make up the WUI, using the Python Perlin Noise library, as mentioned in the 2.4 section (see Fig. 6). These scenarios will be used to train the neural network through several epochs or iterations, where the initial origin of the wildfire defined in the centre of the scenario, allowing the agent to learn to explore the influence of the various variables and decisions in terms of critical wind direction that maximise the propagation towards the WUI and therefore minimise the ASET to reach the WUI through the use of rewards. Following the same procedure, 100 synthetic random scenarios of various dimensions (11 × 11, 15 × 15, and 31 × 31) with a fixed spatial resolution of 300 m were generated to evaluate the performance of the DQN-Agent in scenarios of different and equal

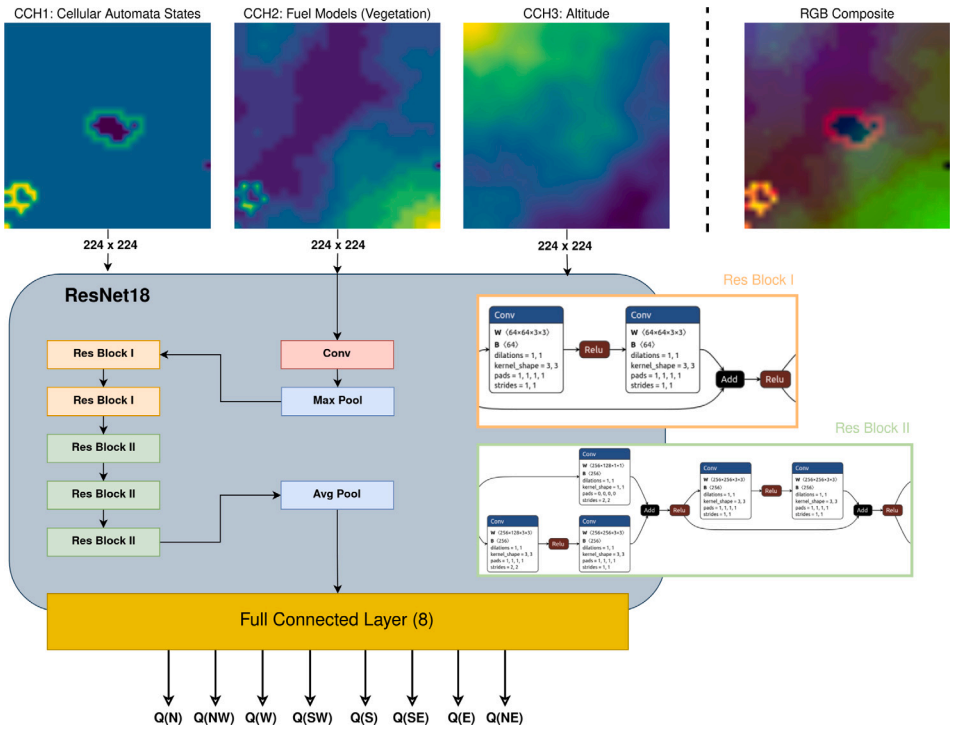

Fig. 5. ResNet18 architecture adaptation for DQN agent.

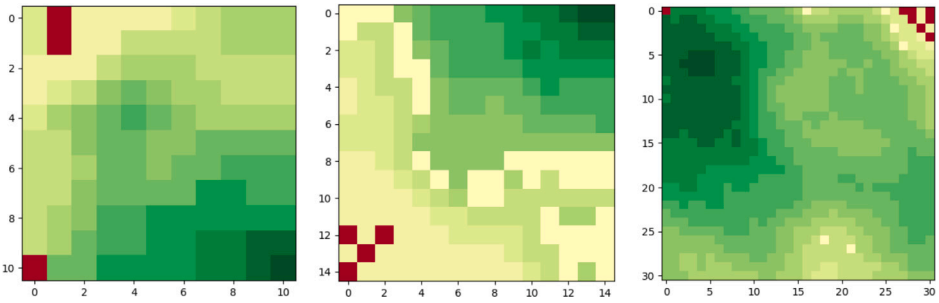

Fig. 6. Synthetic random training and testing scenarios of different dimensions (11 × 11, 15 × 15 and 31 × 31).

Table 3
Rothernmel model verification scenarios.

Scenario	Fuel Model	Slope (%)	Wind (m/s)	Moisture content (%)
S1	Grass (short)	0	0	5
S2	Grass (short)	0	0	20
S3	Grass (short)	0	2	5
S4	Grass (short)	0	2	20
S5	Grass (short)	40	0	5
S6	Grass (short)	40	0	20
S7	Grass (short)	40	2	5
S8	Grass (short)	40	2	20
S9	Chaparral	0	0	5
S10	Chaparral	0	0	20
S11	Chaparral	0	2	5
S12	Chaparral	0	2	20
S13	Chaparral	40	0	5
S14	Chaparral	40	0	20
S15	Chaparral	40	2	5
S16	Chaparral	40	2	20

dimensions to the training scenario, to be compared with the minimum that the stochastic model can achieve in 100 iterations.

Use Case. Finally, by obtaining topographical data of the terrain, vegetation and climatic properties of the area, a real use case was obtained to illustrate the operation of the development and its integration capabilities with GIS systems already deployed (see Figs. 7 and 8). The location of the scenario refers to a real forest fire that started in the locality of El Pinar (Puntagorda) [lon: −17.9776 lat: 28.77111] on July 15, 2023, at 01:20 h in La Palma (Spain) which destroyed 4765 ha and displaced more than 4000 people (Copernicus EMSR671). Owing to the duration of the fire and the human resources involved, this use case serves as an illustration of the performance of the proposed hybrid approach and its capabilities in identifying the MEFS and LFS scenarios, as it could not capture the impact of the various suppression measures carried out by firefighters and other forest response services, and their influence on the spread of the fire. The data collected for the simulation comprised a 10 × 10 km grid, with cells having a spatial resolution of 100 × 100 m. To understand the application of this model to the use case and given that the original ignition point was practically next to the WUI, this has been modified for illustrative purposes. For the characterisation of the various cells that make up the automaton, data were obtained mainly from three sources:

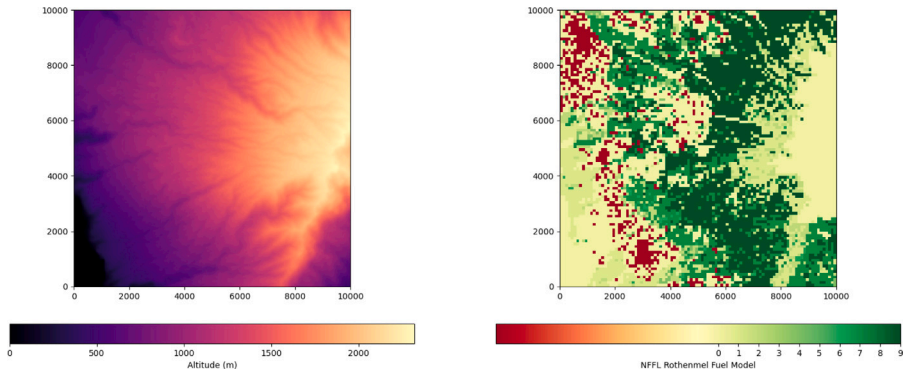

Fig. 7. La Palma use case orography, Rothenmel fuel models and WUI mapping via cellular automata.

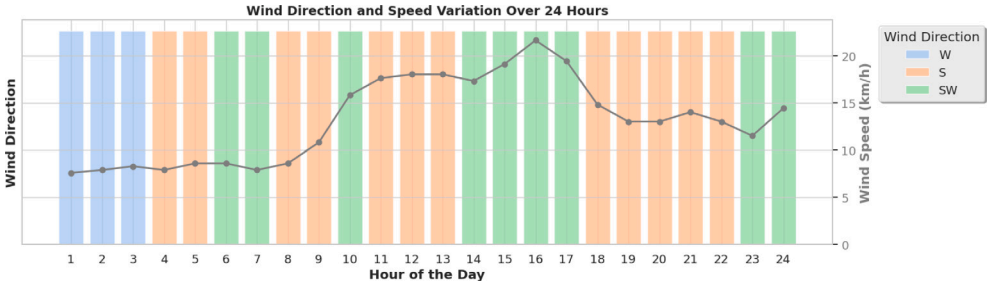

Fig. 8. La Palma (15/07/2023) historical data on wind speed and direction.

- Topographic data: Elevation data were obtained from the Copernicus Global Digital Elevation Model with a spatial resolution of 10 m. European Space Agency and Sinergise (2021)
- Fuel models: Fuel model data were obtained from the European Forest Fire Information System (EFFIS) fuel map (European Forest Fire Information System (EFFIS), 2017). This map was developed as part of the JRC Contract Number 384347, titled "Development of a European Fuel Map". The original mapping identified 42 distinct vegetation complexes, which were subsequently categorised into 10 of the 13 classes defined by the Northern Forest Fire Laboratory (NFFL) fire behaviour models. These models, established by Anderson (1982), provide a standardised classification system for fire behaviour prediction. It is worth noting that the European Fuel Map excludes the last three categories of NFFL models, resulting in a total of 10 fuel classes in the European context.
- Historical wind data: Open data from the Spanish Agencia Estatal de Meteorología (AEMET) (Agencia Estatal de Meteorología, 2024).

4. Results

The initial results verify the correct implementation of the Rothermel model in our development, as shown in Fig. 9. These results demonstrated the influence of various factors collected across different test scenarios. Not only did visual inspection corroborate the proper implementation, but analytical calculations also confirmed the accuracy of the propagation speeds by validating propagation distances. These ranged from cases with limited propagation, such as S1 (101 m) and S2 (71 m), to scenarios that encompassed nearly the entire area and were verified by the time taken to reach the boundaries of the scenario, as observed in S9 (2832 s) and S11 (204 s). As shown in Fig. 10, simulation time decreases significantly with fewer cellular automaton cells, reflecting higher spatial resolution; however, reducing the automaton size (in terms of number of cells) causes a marked increase in relative error in the estimated fire spread area, highlighting a critical trade-off between accuracy and computational efficiency.

Following the model verification, the next step involved training the DQN-Agent using set of training scenarios. During this process, the agent's ability to reach the WUI under various conditions was monitored through two metrics: (1) streak score of number of scenarios solved consecutively in each epoch and (2) relative improvement of ASET with respect to the stochastic model. Fig. 11 presents the evolution of the scores obtained across the different training epochs representing the Maximum Moving (MM) and the Simple Moving Average (SMA) with a Sliding Window (SW) size of 5000, to better understand the evolution of the scores avoiding training noise. Initially, the agent exhibited more chaotic behaviour, partly due to the epsilon value being set to 1 and gradually decreasing to 0 due to the epsilon decay (0.9975) after approximately 1000 epochs, allowing the agent to explore the environment. After this initial phase, the metrics stabilised until the agent reached a point where it could solve a large number of scenarios consecutively with a clear improvement in the time proposed by the stochastic model. It is important to note that the oscillations in the training process also depended heavily on the random nature of the scenarios, which did not obscure the clear improvement trend of the agent. As an important result of this figure, we can conclude that after the training process, the model is able to obtain, in average, better LFS than those proposed by the stochastic model with a relative average improvement of around 20%.

Subsequent to this training process, we proceeded to the use of the synthetic scenarios of different dimensions for the validation of the DQN-Agent, and we could observe an improvement in larger settings than those used for training compared to the minimum value of the stochastic model which is more suitable for smaller scenarios. As can be seen in Fig. 12, DQN-Agent obtains a lower average ASET in certain scenarios and the percentage of times where the DQN model is able to identify a more critical scenario in terms of ASET than the one found by MST in its entire scouting, a trend that increases in larger scenarios, which has some coherence since the exploration required for the stochastic model in larger scenarios should be higher, being necessary to increase the number of iterations.

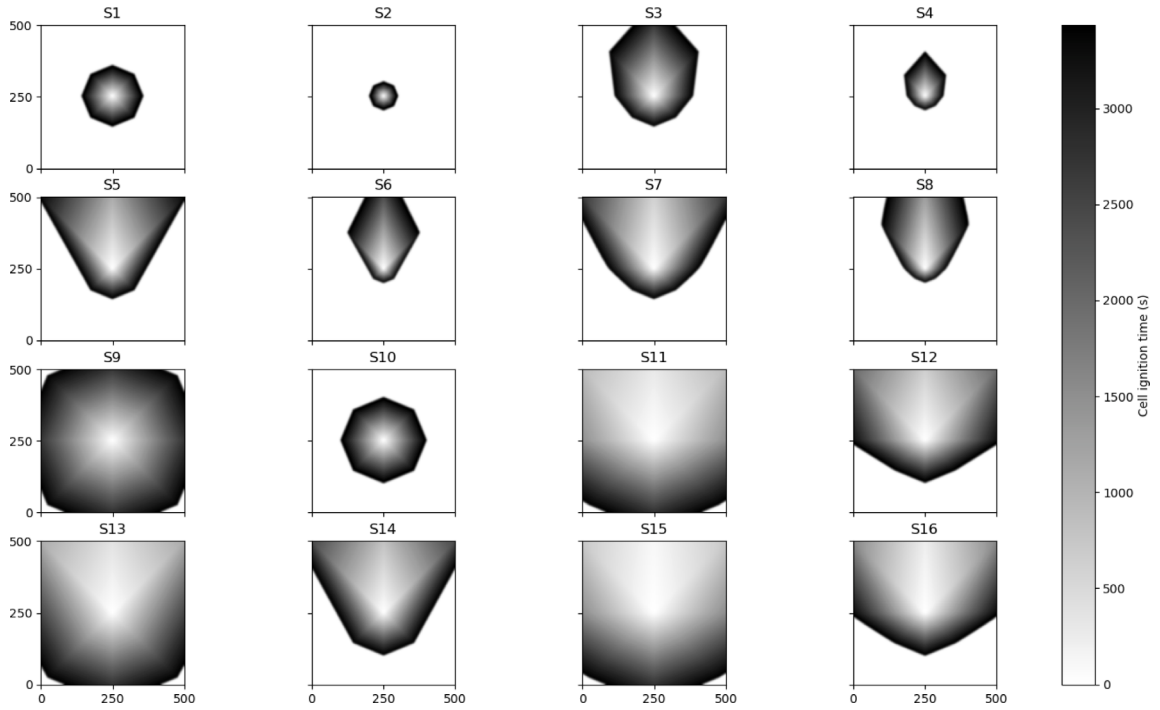

Fig. 9. Visual verification of cell ignition time results of the implementation with the scenarios presented in Table 3.

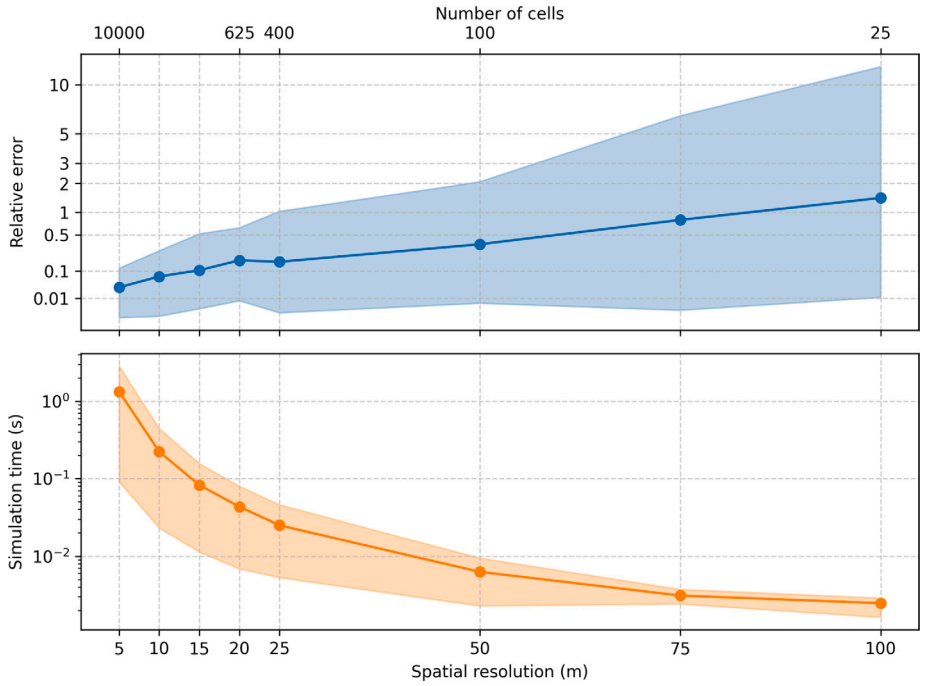

Fig. 10. Relative error (top) showing mean with shaded max–min range; average computational time (bottom) with mean line and min–max band across scenarios presented in Table 3.

On the other hand, Fig. 13 shows a comparison of the simulation times for each of the approaches using a commercial computer without GPU with a Ryzen 5700 CPU, so that both models are on equal footing. With these results, it is clear that the simulation time of the DQN-Agent is significantly shorter once trained, as it is only necessary to carry out one prediction, compared to the time of one iteration of the stochastic model, where multiple iterations are required to reach the LFS. It should also be considered that the evolution in terms of scenario size in the case of the stochastic model is much larger compared to the DQN-Agent. It is true that these times, in the stochastic case, could be reduced by parallelisation techniques or by seeking convergence by other methods instead of seeking a fixed number of iterations.

The figures of the last results presented correspond to the application of the complete system to the La Palma use case, allowing us to see how it works and the integration of external data sources. The graphs in Fig. 14 show the functionalities that are provided. The first piece of information provided by the system is based on the stochastic estimation of the risk in the area in terms of cell range probability given

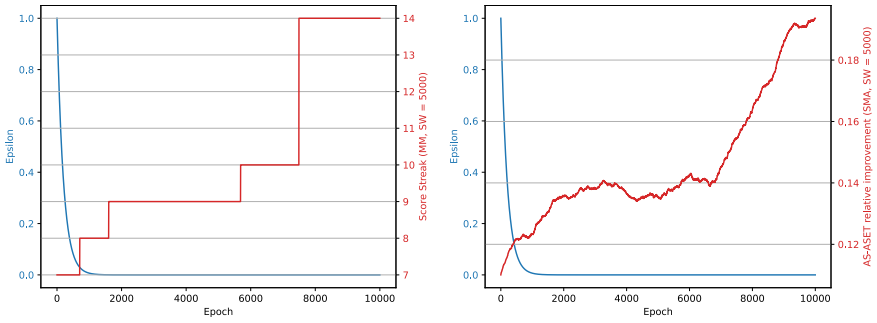

Fig. 11. Results of the metrics assessed with the MM (left) and the SMA (right) during the DQN-Agent training process.

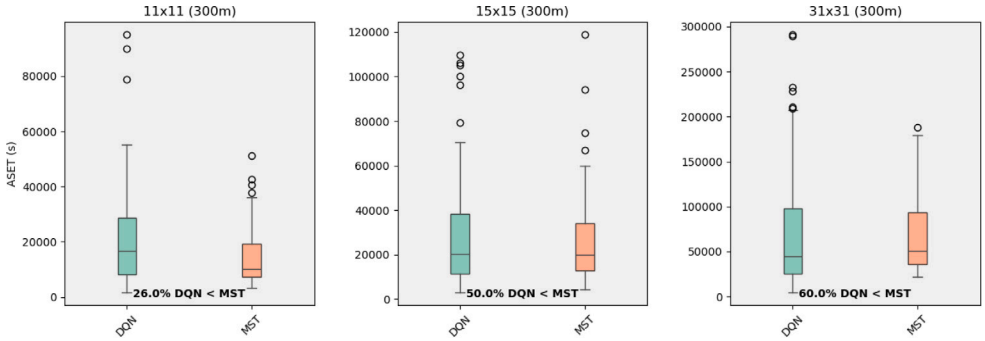

Fig. 12. Minimum obtained ASET comparison between stochastic model (MST) and DQN-Agent(DQN) and percentage of scenarios where DQN provides ASETs lower than those estimated by stochastic model.

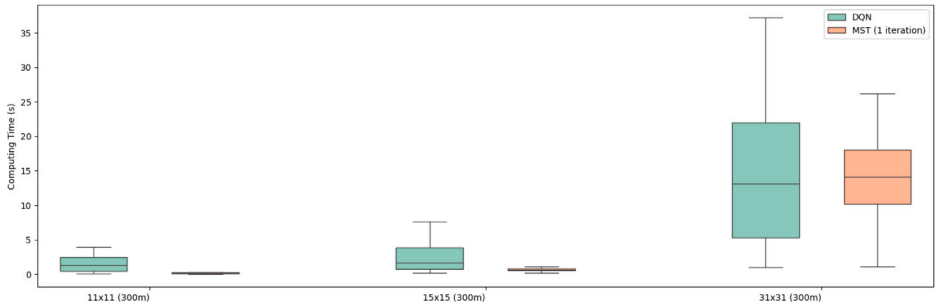

Fig. 13. Computing time comparison between stochastic model (MST) and DQN-Agent(DQN).

a fixed fire ignition origin (marked in yellow in the figure) and the initial distribution of meteorological conditions presented in Fig. 8.

In the same figure, the MEFS is presented for the specific use case using the same ignition source and meteorological conditions of the event detailed above, which finally provide an estimated ASET metric from the origin to the WUI. Finally, to test the capabilities of the DQN-Agent, the same fire origin has been used to understand what would be the LFS in those conditions, where the wind is not determined by a probability distribution and what is sought is to find the worst-case scenario. These results are shown in the same figure that presents the final propagation LFS associated to the critical wind direction profile in that case related to a NW wind profile. From this figure, it can be seen that the LFS presents an ASET significantly lower than the one found by the MEFS in an acceptable computation time to be considered real-time.

5. Discussion & conclusions

This study presents an innovative computational framework that integrates cellular automata modelling of wildfire propagation with stochastic and deep reinforcement learning techniques to address critical challenges in wildfire surface propagation prediction for evacuation and response management. This approach combines a custom Rothermel surface fire spread model implemented via cellular automata with Deep Q-Networks (DQN) and stochastic modelling to predict wildfire spread probability distributions and identify limiting fire scenarios. The methodology was verified and trained using synthetic scenarios and was demonstrated in a real-world case. The integration of stochastic simulation and reinforcement learning techniques offers several key advantages and applications. Compared to traditional deterministic wildfire spread models (Finney, 1998; Lautenberger, 2013; Vasconcelos & Guertin, 1992), which rely on predefined input parameters and static propagation rules, the proposed framework demonstrates superior adaptability and scenario-specific responsiveness. These classical models, while robust in structured environments, lack the capacity to autonomously identify critical fire trajectories under uncertainty, a capability that is central to the DQN-Agent's performance. The results obtained from the simulation experiments support the effectiveness of this hybrid approach. The DQN-Agent consistently outperformed the stochastic baseline in identifying LFS, achieving an average ASET reduction of approximately 20% across synthetic test scenarios. This improvement was particularly evident in larger grid configurations (31 × 31), where the agent demonstrated superior generalisation and

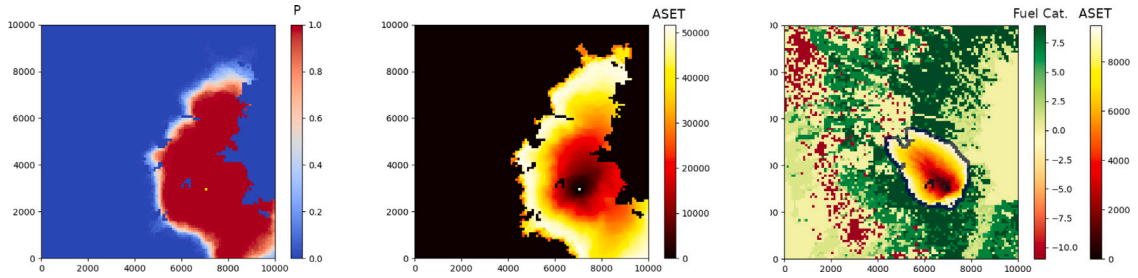

Fig. 14. Fire extent risk probability map (left), ASET of the stochastically estimated MEFS scenario depending on the fire origin (centre) and LFS identified by the DQN-Agent (right) for the La Palma case study.

efficiency. Moreover, the agent's performance remained robust across randomised terrain, fuel mosaics, and wind conditions, indicating resilience to overfitting. The La Palma case study further validated the model's practical applicability, where the DQN-Agent identified a more critical LFS than the MEFS baseline under real topographic and meteorological conditions. These findings highlight the model's potential for real-time deployment in WUI risk assessment and evacuation planning.

This probabilistic approach allows for a more accurate risk assessment and evacuation planning in WUI areas compared to traditional deterministic models (Asensio et al., 2021; Finney, 1998; Lautenberger, 2013; Vasconcelos & Guertin, 1992), capturing the inherent unpredictability of wildfires but reducing the computational load. The use of cellular automata provides seamless integration with GIS data and digital twins, improving the model's spatial resolution and real-world applicability, a very interesting capability that has already been demonstrated by other studies (Green et al., 1995) and is very interesting for establishing synergies with existing evacuation models (González-Villa et al., 2022; Wahlqvist et al., 2021). Additionally, the combination of stochastic modelling with DQN-Agent ability to identify LFS offers valuable insights for proactive wildfire management and mitigation strategies. However, this approach has limitations that warrant further investigation. Despite its strengths, the proposed model should be benchmarked against probabilistic Monte Carlo-based approaches (Boychuk et al., 2009; Waeselynck & Saah, 2024), which offer high-fidelity uncertainty quantification but at a significantly higher computational cost. In contrast, the DQN-Agent achieves comparable or better scenario identification in over 60% of evaluated cases (see Fig. 12), with simulation times reduced by an order of magnitude (see Fig. 13), making it more suitable for real-time decision support. The custom Rothermel model used in this study omits certain complex fire behaviours such as canopy spread, spotting, and smouldering (Albini, 1979; Albini et al., 2012; Purnomo et al., 2024), which could be considered in some WUI scenarios. The exclusion of crown fires, spotting, and smouldering is justified by the dominance of surface fire spread dynamics in the WUI settings addressed, where fine fuels and managed fuel beds reduce canopy-fire prevalence. These simplifications enable real-time prediction and prioritise computational tractability. However, scenarios with dense continuous canopy or extreme winds are acknowledged as underrepresented. Future work will extend the CA framework to incorporate layers modelling crown fire transition, ember spotting, and subcanopy smouldering, increasing the model's applicability to broader WUI and wildland environments.

Through the results of application La Palma use case, it is clear that more sophisticated fire spread models are required. These should involve multiple layers of information interacting with each other to address these limitations, including the ability to consider some real-time extinction measures. The model currently excludes dynamic firefighting interventions such as suppression lines or aerial retardant drops, which limits its immediate utility for emergency response as suppression effects on ASET. However, this tool is primarily conceived as a predictive decision support system. Identifies critical fire spread scenarios, which can be used to guide the planning of mitigation measures, as has been seen in other studies (Granda Chico, 2024). Preventive strategies, such as firebreaks, can be modelled as automaton cells without fuel (model 0 in Table 1). This configuration represents physical barriers within the fuel model layer of the cellular automaton, as illustrated in Fig. 2. Similarly, suppression interventions, such as aerial retardant drops or other firefighting actions, can be represented as localised increases in fuel moisture within the combustion framework described above. These effects can be incorporated as additional layers in both the automaton and DQN-Agent architecture (see Fig. 5), allowing environment perception to adapt after the implementation of such preventive or response measures. Integrating real-time suppression modelling into the cellular automata and reinforcement learning framework remains a key priority to simulate intervention effects accurately and enhance operational planning. Additionally, ongoing efforts aim to extend the model's use beyond planning to emergency response and post-event resilience analysis, facilitating ecosystem recovery assessments and enhancing future wildfire management strategies. Another important aspect to consider is the inclusion of internal fuel moisture content (Rothermel, 1986) in the model instead of a homogeneous value for the entire area, which is a challenge because of the scarcity of studies that directly correlate in a generalised way this parameter with the various fuel categories and the external ambient humidity and temperature, which can be obtained from meteorological systems. Furthermore, future work could also consider the improvement of models including the opportunity to train the agent with real data obtained by GIS in an automated manner instead of synthetic scenarios, always following the same methodology outlined above. The performance of DQN-Agent could potentially be enhanced through more advanced reinforcement learning architectures such as Double DQN (Van Hasselt, Guez, & Silver, 2016) or Dueling DQN (Wang et al., 2016). The study of alternative exploration strategies, such as parameter noise (Plappert et al., 2017) or intrinsic motivation (Pathak, Agrawal, Efros, & Darrell, 2017) could also improve the agent's ability to identify critical scenarios. Although the stochastic simulation approach is valuable for capturing uncertainty, it can be computationally intensive in some scenarios, as presented in Figs. 10 and 13. It is worth noting that this process is inherently parallelizable, and future work should explore optimisation techniques to reduce simulation times, particularly for complex scenarios such as potential megafires. Although the training approach was specifically designed to minimise overfitting — using thousands of synthetically generated scenarios with randomised landscapes, fuel mosaics, and weather conditions to prevent adaptation to a single geography or event type — its evaluation has so far been limited. Current assessments include only synthetic environments and a single real-world case: the La Palma scenario, representative of Mediterranean woodland. As a result, the model's ability to generalise to other ecosystem types (such as boreal forests or shrublands) remains unverified. Future work will focus on adapting the model and conducting rigorous validation using historical fire data from diverse biomes. This will involve systematic comparisons of transferability and refinement of input feature selection to ensure robust performance across multiple ecosystems.

One promising avenue for optimisation is the implementation of dynamic spatial resolutions for the cellular automaton simulation. This

approach would allow for higher resolution in areas of immediate concern or rapid fire spread, while using lower resolution in less critical areas. Such an adaptive resolution strategy can significantly improve computational efficiency without sacrificing accuracy in critical regions. The model's deployment in real wildfire emergencies benefits from a response window that is realistically aligned with operational needs, thus not requiring intensive computational optimisation. Studies show that current predictive wildfire systems provide actionable forecasts within minutes to hours, ensuring timely decision support without excessive computational burden (Jain et al., 2020b). This balance between computational demand and response time supports effective emergency management. Decision support systems for forest fire management are most effective with data update intervals under 30 min, ideally between 5 and 30 min, to ensure timely and actionable decisions in rapidly evolving fire incidents (Granda Chico, 2024). Achieving this dynamic approach requires careful calibration to quantify the degree of precision loss and establish an optimal balance between computational efficiency and predictive fidelity. This fine-tuning and benchmarking will be the focus of forthcoming studies. Additionally, future research will explore advanced surrogate and reduced-order models, hardware acceleration, and parallelisation strategies to enhance scalability further. It is important to emphasise that this study represents an initial exploratory step to assess the capabilities of integrating cellular automata, stochastic simulation, and deep reinforcement learning for wildfire prediction and management. The results demonstrate the potential of this approach; however, extensive validation against real-world wildfire data and comparison with existing operational models are necessary to fully assess the practical utility and limitations of the framework.

CRediT authorship contribution statement

Javier González-Villa: Conceptualization, Formal analysis, Methodology, Software, Writing – original draft, Validation. **David Lázaro:** Conceptualization, Formal analysis, Supervision, Writing – review & editing, Validation. **Arturo Cuesta:** Investigation, Visualization, Data curation. **Adriana Balboa:** Investigation, Resources, Data curation, Validation. **Daniel Alvear:** Supervision, Project administration, Funding acquisition. **Mariano Lázaro:** Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This publication is derived from the projects 'Real time Application for Protecting People In Disasters (RAPPID)', Grant TED2021-132410B-I00 funded by MICIU/AEI/10.13039/501100011033 and by "European Union NextGenerationEU/PRTR" and project "Integrated Management for Prevention, Extinction and Reforestation of Forest Fires (GAIA)" Grant PLEC2023-010303 funded by MICIU/AEI/10.13039/501100011033.

Data availability

The data that support the findings of this study are openly available at the following URL:

https://zenodo.org/records/14991404.

References

Agencia Estatal de Meteorología (2024). AEMET OpenData: Open data catalogue. https://www.aemet.es/en/datos_abiertos/catalogo. (Accessed 6 September 2024).

Albini, F. A. (1976). *Computer-based models of wildland fire behavior: a user's manual*. Intermountain Forest and Range Experiment Station, Forest Service, US Department of Agriculture.

Albini, F. A. (1979). *Spot fire distance from burning trees: a predictive model: vol. 56*, Intermountain Forest and Range Experiment Station, Forest Service, US Department of Agriculture.

Albini, F. A., Alexander, M. E., & Cruz, M. G. (2012). A mathematical model for predicting the maximum potential spotting distance from a crown fire. *International Journal of Wildland Fire*, *21*(5), 609–627. http://dx.doi.org/10.1071/WF11020.

Anderson, H. E. (1983). Predicting wind-driven wild land fire size and shape. http://dx.doi.org/10.5962/bhl.title.69035.

Andrews, P. L. (2012). *266*, *Modeling wind adjustment factor and midflame wind speed for Rothermel's surface fire spread model: Gen. Tech. Rep. RMRS-GTR-266*, (p. 39). Fort Collins, CO: US Department of Agriculture, Forest Service, Rocky Mountain Research Station, http://dx.doi.org/10.2737/rmrs-gtr-266.

Andrews, P. L., Cruz, M. G., & Rothermel, R. C. (2013). Examination of the wind speed limit function in the rothermel surface fire spread model. *International Journal of Wildland Fire*, *22*(7), 959–969. http://dx.doi.org/10.1071/WF12122.

Andrews, P. L., et al. (2018). The rothermel surface fire spread model and associated developments: A comprehensive explanation. http://dx.doi.org/10.2737/RMRS-GTR-371.

Asensio, M., Ferragut, L., Álvarez, D., Laiz, P., Cascón, J., Prieto, D., et al. (2021). Phyfire: an online gis-integrated wildfire spread simulation tool based on a semiphysical model. In *Applied mathematics for environmental problems* (pp. 1–20). Springer, http://dx.doi.org/10.1007/978-3-030-61795-0_1.

Association, N. F. P., et al. (2006). *Performance-Based Standard for Fire Protection for Light Water Reactor Electric Generating Plants*. National Fire Protection Association.

Bañuelos-Ruedas, F., Angeles-Camacho, C., & Rios-Marcuello, S. (2010). Analysis and validation of the methodology used in the extrapolation of wind speed data at different heights. *Renewable and Sustainable Energy Reviews*, *14*(8), 2383–2391. http://dx.doi.org/10.1016/j.rser.2010.05.001.

Beneduci, R., & Mascali, G. (2024). Forest fire spreading: A nonlinear stochastic model continuous in space and time. *Studies in Applied Mathematics*, *153*(1), Article e12696. http://dx.doi.org/10.1111/sapm.12696.

Beyki, S. M., Santiago, A., Laím, L., & Craveiro, H. D. (2023). Evacuation simulation under threat of wildfire—an overview of research, development, and knowledge gaps. *Applied Sciences*, *13*(17), 9587. http://dx.doi.org/10.3390/app13179587.

Boychuk, D., Braun, W. J., Kulperger, R. J., Krougly, Z. L., & Stanford, D. A. (2009). A stochastic forest fire growth model. *Environmental and Ecological Statistics*, *16*, 133–151. http://dx.doi.org/10.1007/s10651-007-0079-z.

Buch, J., Williams, A. P., Juang, C. S., Hansen, W. D., & Gentine, P. (2023). Smlfire1.0: a stochastic machine learning (SML) model for wildfire activity in the western United States. *Geoscientific Model Development*, *16*(12), 3407–3433. http://dx.doi.org/10.5194/gmd-16-3407-2023.

Cruz, M. G., Alexander, M. E., & Wakimoto, R. H. (2004). Modeling the likelihood of crown fire occurrence in conifer forest stands. *Forest Science*, *50*(5), 640–658. http://dx.doi.org/10.1093/forestscience/50.5.640.

Danta, J. M., Egorova, V. N., & Pagnini, G. (2024). Predicting the arrival of the unpredictable: An approach for foreseeing the transition to chaos of wildfire propagation. *Communications in Nonlinear Science and Numerical Simulation*, *138*, Article 108190. http://dx.doi.org/10.1016/j.cnsns.2024.108190.

Etherington, T. R. (2022). Perlin noise as a hierarchical neutral landscape model. *Web Ecology*, *22*(1), 1–6. http://dx.doi.org/10.5194/we-22-1-2022.

European Forest Fire Information System (EFFIS) (2017). European fuel map. Based on JRC Contract Number 384347 on the "Development of a European Fuel Map".

European Space Agency, & Sinergise (2021). Copernicus global digital elevation model. http://dx.doi.org/10.5069/G9028PQB, Accessed: 2024-07-03.

Finney, M. A. (1998). *FARSITE, Fire Area Simulator–model development and evaluation*. (4), US Department of Agriculture, Forest Service, Rocky Mountain Research Station, http://dx.doi.org/10.2737/RMRS-RP-4.

Frandsen, W. H. (1971). Fire spread through porous fuels from the conservation of energy. *Combustion and Flame*, *16*(1), 9–16.

Freire, J. G., & DaCamara, C. C. (2019). Using cellular automata to simulate wildfire propagation and to assist in fire management. *Natural Hazards and Earth System Sciences*, *19*(1), 169–179. http://dx.doi.org/10.5194/nhess-19-169-2019.

Ganapathi Subramanian, S., & Crowley, M. (2018). Using spatial reinforcement learning to build forest wildfire dynamics models from satellite images. *Frontiers in ICT*, *5*, 6. http://dx.doi.org/10.3389/fict.2018.00006.

González-Villa, J., Cuesta, A., Alvear, D., & Balboa, A. (2022). Evacuation management system for major disasters. *Applied Sciences*, *12*(15), 7876. http://dx.doi.org/10.3390/app12157876.

Granda Chico, B. (2024). Investigación operativa para la extinción de incendios forestales: El "team orienteering problem" con ventanas de tiempo variables.

Green, M. E., DeLuca, T. F., & Kaiser, K. W. (2020). Modeling wildfire using evolutionary cellular automata. In *Proceedings of the 2020 genetic and evolutionary computation conference* (pp. 1089–1097). http://dx.doi.org/10.1145/3377930.3389836.

Green, K., Finney, M., Campbell, J., Weinstein, D., & Landrum, V. (1995). Fire! using GIS to predict fire behavior. *Journal of Forestry*, *93*(5), 21–25. http://dx.doi.org/10.1093/jof/93.5.21.

Gwynne, S. M., Ronchi, E., Wahlqvist, J., Cuesta, A., González-Villa, J., Kuligowski, E. D., et al. (2023). Roxborough park community wildfire evacuation drill: Data collection and model benchmarking. *Fire Technology*, *59*(2), 879–901. http://dx.doi.org/10.1007/s10694-023-01371-1.

Hajian, M., Melachrinoudis, E., & Kubat, P. (2016). Modeling wildfire propagation with the stochastic shortest path: A fast simulation approach. *Environmental Modelling & Software*, *82*, 73–88. http://dx.doi.org/10.1016/j.envsoft.2016.03.012.

Hurley, M. J., Gottuk, D. T., Hall, Jr., J. R., Harada, K., Kuligowski, E. D., Puchovsky, M., et al. (2015). *SFPE handbook of fire protection engineering*. Springer, http://dx.doi.org/10.1007/978-1-4939-2565-0.

Jain, P., Coogan, S. C., Subramanian, S. G., Crowley, M., Taylor, S., & Flannigan, M. D. (2020a). A review of machine learning applications in wildfire science and management. *Environmental Reviews*, *28*(4), 478–505.

Jain, P., Coogan, S. C., Subramanian, S. G., Crowley, M., Taylor, S., & Flannigan, M. D. (2020b). A review of machine learning applications in wildfire science and management. *Environmental Reviews*, *28*(4), 478–505. http://dx.doi.org/10.1139/er-2020-0019.

Lautenberger, C. (2013). Wildland fire modeling with an Eulerian level set method and automated calibration. *Fire Safety Journal*, *62*, 289–298. http://dx.doi.org/10.1016/j.firesaf.2013.08.014.

Leblon, B., Bourgeau-Chavez, L., & San-Miguel-Ayanz, J. (2012). Use of remote sensing in wildfire management. *Sustainable Development-Authoritative and Leading Edge Content for Environmental Management*, 55–82. http://dx.doi.org/10.5772/45829.

Lier, M., Köhl, M., Korhonen, K. T., Linser, S., Prins, K., & Talarczyk, A. (2022). The new EU forest strategy for 2030: a new understanding of sustainable forest management? *Forests*, *13*(2), 245. http://dx.doi.org/10.3390/f13020245.

Linn, R. R., & Harlow, F. H. (1997). *FIRETEC: a transport description of wildfire behavior*: *Tech. Rep.*, Los Alamos National Lab.(LANL), Los Alamos, NM (United States).

Liu, N., Lei, J., Gao, W., Chen, H., & Xie, X. (2021). Combustion dynamics of large-scale wildfires. *Proceedings of the Combustion Institute*, *38*(1), 157–198. http://dx.doi.org/10.1016/j.proci.2020.11.006.

McArthur, A. G. (2000). *The Tasmanian bushfires of 7th February 1967 and associated fire behaviour characteristics*.

Mell, W. (2010). User guide to WFDS-this is a work in progress. *NIST, US*, *21*.

MITECO (2023). *Evaluación y zonificación del riesgo en la interfaz urbano-forestal en España*: *Tech. Rep.*, Ministerio para la Transición Ecológica y el Reto Demográfico, URL https://www.miteco.gob.es/content/dam/miteco/es/biodiversidad/temas/incendios-forestales/cap_1_evaluacion_riesgo_iuf_tcm30-278867.pdf. (Accessed October 2025).

Morvan, D. (2011). Physical phenomena and length scales governing the behaviour of wildfires: a case for physical modelling. *Fire Technology*, *47*(2), 437–460. http://dx.doi.org/10.1007/s10694-010-0160-2.

Murray, L., Castillo, T., Carrasco, J., Weintraub, A., Weber, R., de Diego, I. M., et al. (2024). Advancing forest fire prevention: Deep reinforcement learning for effective firebreak placement. arXiv preprint arXiv:2404.08523.

Nguyen, C., Schlesinger, K. J., Han, F., Gür, I., & Carlson, J. M. (2019). Modeling individual and group evacuation decisions during wildfires. *Fire Technology*, *55*, 517–545. http://dx.doi.org/10.1007/s10694-018-0770-7.

Olmo, J. M., Ruiz, J., Planelles, R., & Hernando, C. (2013). Characterization of wildland-urban interfaces for fire prevention in the province of Valencia (Spain). *Forest Systems*, *22*(2), 249–254. http://dx.doi.org/10.5424/fs/2013222-03985.

Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning* (pp. 2778–2787). PMLR, http://dx.doi.org/10.1109/cvprw.2017.70.

Plappert, M., Houthooft, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., et al. (2017). Parameter space noise for exploration. http://dx.doi.org/10.48550/arXiv.1706.01905, arXiv preprint arXiv:1706.01905.

Purnomo, D. M., Christensen, E. G., Fernandez-Anez, N., & Rein, G. (2024). BARA: cellular automata simulation of multidimensional smouldering in peat with horizontally varying moisture contents. *International Journal of Wildland Fire*, *33*(2), NULL–NULL. http://dx.doi.org/10.1071/wf23042.

Ramírez, J. A. H., Lendechy, C. R. A., Castillo, W. A., Toscano, G. M. O. C., Medina, A. P. L., Vargas, R., et al. (2016). *Ciencias ambientales*. Universidad de Xalapa, A. C..

Rothermel, R. C. (1972). *A mathematical model for predicting fire spread in wildland fuels*: *vol. 115*, Intermountain Forest & Range Experiment Station, Forest Service, US ....

Rothermel, R. C. (1986). *Modeling moisture content of fine dead wildland fuels: input to the BEHAVE fire prediction system*. (359), US Department of Agriculture, Forest Service, Intermountain Research Station, http://dx.doi.org/10.2737/int-rp-359.

Rothermel, R. C., & Anderson, H. E. (1966). *Fire spread characteristics determined in the laboratory*: *vol. 30*, Intermountain Forest & Range Experiment Station, Forest Service, US ....

Singh, H., Ang, L.-M., Lewis, T., Paudyal, D., Acuna, M., Srivastava, P. K., et al. (2024). Trending and emerging prospects of physics-based and ML-based wildfire spread models: A comprehensive review. *Journal of Forestry Research*, *35*(1), 135. http://dx.doi.org/10.1007/s11676-024-01783-x.

Sun, X., Li, N., Chen, D., Chen, G., Sun, C., Shi, M., et al. (2024). A forest fire prediction model based on cellular automata and machine learning. *IEEe Access*, http://dx.doi.org/10.70127/irjedt.vol.8.issue05.734.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press, http://dx.doi.org/10.1109/TNN.1998.712192.

Thompson, M. P., & Calkin, D. E. (2011). Uncertainty and risk in wildland fire management: a review. *Journal of Environmental Management*, *92*(8), 1895–1909. http://dx.doi.org/10.1016/j.jenvman.2011.03.015.

Van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*: *vol. 30*, (1), http://dx.doi.org/10.48550/arXiv.1509.06461.

Vasconcelos, M. J., & Guertin, D. P. (1992). FIREMAP-simulation of fire growth with a geographic information system. *International Journal of Wildland Fire*, *2*(2), 87–96. http://dx.doi.org/10.1071/WF9920087.

Velasquez, W., Munoz-Arcentales, A., Bohnert, T. M., & Salvachúa, J. (2019). Wildfire propagation simulation tool using cellular automata and GIS. In *2019 international symposium on networks, computers and communications* (pp. 1–7). IEEE, http://dx.doi.org/10.1109/ISNCC.2019.8909129.

Waeselynck, V., & Saah, D. (2024). Importance sampling for cost-optimized estimation of burn probability maps in wildfire Monte Carlo simulations. *Fire*, *7*(12), 455. http://dx.doi.org/10.20944/preprints202410.0879.v1.

Wahlqvist, J., Ronchi, E., Gwynne, S. M., Kinateder, M., Rein, G., Mitchell, H., et al. (2021). The simulation of wildland-urban interface fire evacuation: The WUINITY platform. *Safety Science*, *136*, Article 105145. http://dx.doi.org/10.1016/j.ssci.2020.105145.

Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., & Freitas, N. (2016). Dueling network architectures for deep reinforcement learning. In *International conference on machine learning* (pp. 1995–2003). PMLR.