

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Data-driven Applications and Decision Making Models in Natural Resources

Permalink

<https://escholarship.org/uc/item/1pz7c83m>

Author

Pais, Cristobal

Publication Date

2021

Peer reviewed|Thesis/dissertation

Data-driven Applications and Decision Making Models in Natural Resources

by

Cristobal Andres Pais Martinez

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Zuo-Jun Max Shen, Chair

Professor Marta Gonzalez

Professor Paul Grigas

Professor Zeyu Zheng

Spring 2021

Data-driven Applications and Decision Making Models in Natural Resources

Copyright 2021
by
Cristobal Andres Pais Martinez

Abstract

Data-driven Applications and Decision Making Models in Natural Resources

by

Cristobal Andres Pais Martinez

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Professor Zuo-Jun Max Shen, Chair

The destructive potential of wildfires has been exacerbated by climate change, causing their frequencies and intensities to continuously increase globally. In this context, increasing wildfire activity across the globe has become an urgent issue with enormous ecological and social impacts [61, 176]. Wildfires have consumed important areas and forest resources, as a result, fire management expenditures have increased and thousands of homes and many lives have been lost [56, 83, 189, 272]. Moreover, they have significantly impacted biodiversity and greenhouse gas emissions at a global scale [24, 205, 237].

The current incidents across the globe highlight the need for preemptive policy measures to reduce the risk of fire occurrence [102], managing the land in an effective way to protect natural forests, agricultural areas, and human lives [73]. These concepts are included in what is known as FireSmart Forest Management (FSFM) [340]. This paradigm considers opportunities in three dimensions: i) decrease of the fire behavior potential of the landscape, ii) reduction of the potential for fire ignitions, and iii) increase in the fire suppression capability.

This dissertation aims at advancing the theory, practice, and large-scale implementation of complex data-driven decision making and machine learning models in the context of landscape management under wildfire risk, integrating Operations Research, Computer Science, and Data Science techniques. We focus our efforts on the understanding, evaluation, and development of effective prevention and mitigation policies, with the potential of being implemented practice, as well as exploring and developing new FSFM techniques.

We divided our study into three main aspects: Simulation, Decision-Making, and Machine Learning. In Chapter 1, we focus on the development and evaluation of an accurate, flexible, and efficient wildfire simulation model that can be integrated with data-driven decision-making models. Empirical results on thousands of simulations show the high performance of the model compared to existing solutions, highlighting its accuracy with real-life instances. We then focus our efforts on its generalization in Chapter 2, seeking to adjust its main

parameters to mimic the fire spread behavior observed in different regions of the world where no empirical models are available. This, exploiting historical information for training purposes using derivative-free optimization techniques to adjust the parameters of the model, allowing us to capture current wildfire dynamics. Experiments performed on datasets located in different regions of the world show the potential of the proposed method.

Second, in Chapters 3 and 4, we explore the integration of this model with landscape planning decision-making models to derive robust fuel treatment policies to mitigate expected losses due to future wildfire events, generating fire resilient landscapes. We study and compare complex network algorithms to develop a mathematical model denoted Downstream Protection Value, capturing the importance of the different components of the land to provide a natural prioritization of where mitigation actions should be implemented. An optimization framework incorporating multiple variables to analyze the inherent trade-offs involved in the planning process is developed, providing practitioners and researchers with an open-source decision support system implementation involving multiple and potentially, opposing objectives. We evaluate the performance of the proposed mathematical model compared to existing solutions, highlighting its superior performance with thousands of experiments involving uncertainty for landscapes located in North America. Several extensions are discussed providing future research directions in the field. We further expand this framework in Chapter 5, incorporating wildfire suppression strategies derived from a novel multi-agent decision-making model. In this application, a group of agents is deployed to the field once an ignition or active fire is detected with the aim of containing or stopping it as soon as possible. The sequential and temporal dimensions of the problem become a challenge to apply traditional modeling techniques. We develop a deep reinforcement learning algorithm focused on exploiting the collaboration and coordination between independent agents. Extensive computational results demonstrate the impact of including local rewards and the concept of sub-groups of agents in the context of centralized training and decentralized execution algorithms, leading to more complex and effective collaboration strategies between agents belonging to the same group and the environment in general.

As part of our decision support system expansion, we build end-to-end machine learning models to understand the wildfire phenomenon from a large-scale perspective. For this, in Chapter 6, we explore the impact of different landscape compositions on the risk of wildfire ignition using remote sensing data to support challenging landscape planning decisions. Using a custom convolutional neural network model integrated with state-of-the-art visualization techniques, we highlight the main areas of interest for the deep learning model, focusing our efforts on the interpretation of the results. This, to open the artificial intelligence black-box to fully understand the rationale behind the results and the different risk levels associated with characteristic spatial patterns observed in the land. We validate our results with previous studies using similar datasets, noting how the proposed model significantly surpasses their predictive performance.

Finally, in Chapter 7, we develop a global study of the main characteristics and drivers of

wildfire regimes consolidating observations for almost two decades of wildfires. We classify and delineate regions of the world sharing similar fire activity as well as identifying their driving factors to support national or regional wildfire prevention/mitigation policies using a variety of machine learning techniques. To the best of our knowledge, this represents the first study that defines fire regimes spatially at a global scale bridging existing knowledge gaps between global and regional fire studies.

Our results represent an attempt at improving the integration of multiple disciplines in the context of effective data-driven decision making under natural hazards uncertainty. We hope that this research can serve as a motivation to expand the field's perspective with high-impact applied projects involving mathematical, ecological, economic, data, and social sciences.

To all my friends and family: my life is great everyday because of you

Contents

Contents	ii
List of Figures	v
List of Tables	xvii
1 Modeling and simulating wildfires	1
1.1 Introduction	1
1.2 Background	3
1.3 Cell2Fire growth simulator	5
1.4 Comparison Methodology: Fire growth validation	10
1.5 Wildfire Management	12
1.6 Experimental Instances	17
1.7 Computational Implementation	19
1.8 Results and Discussion	20
1.9 Conclusions	30
2 Learning from historical data: a general fire spread model using derivative-free optimization	31
2.1 Introduction	31
2.2 A General Fire Spread Model	34
2.3 Data Mining: Historical Fire Scars	43
2.4 Learning Propagation Patterns	45
2.5 Derivative-free optimization	47
2.6 Experiments	50
2.7 Fire spread: learning & adjustments	52
2.8 Conclusions	56
3 Decision making under wildfire uncertainty: A data-driven optimization approach for landscape management	60
3.1 Introduction	60
3.2 Methods	63

3.3	Results and Discussion	77
3.4	Conclusions	84
4	Decision making under wildfire uncertainty: a multi-criteria extension	86
4.1	Introduction	86
4.2	Data Extraction and Processing	88
4.3	Wildfire simulation	91
4.4	Risk analysis	92
4.5	Case study areas	96
4.6	Experiments	98
4.7	Results and Discussion	100
4.8	Conclusions	106
5	Wildfire suppression: A coordinated multi-agent approach	108
5.1	Introduction	108
5.2	Theoretical Background: Deep Reinforcement Learning (DRL)	110
5.3	A Wildfire Suppression Environment	111
5.4	Algorithms implemented	113
5.5	A local reward extension of COMA: SubG-COMA	117
5.6	Experiments	118
5.7	Results & Discussion	119
5.8	Conclusions	123
6	Landscape planning and wildfire risk: opening the deep learning black box.	125
6.1	Introduction	125
6.2	Data mining, processing, and inputs	128
6.3	Deep learning framework	131
6.4	Computational implementation	134
6.5	Results and Discussion	134
6.6	Conclusions	147
7	Wildfires in a nutshell: a global-scale study	149
7.1	Introduction	149
7.2	Global Wildfire Datasets	153
7.3	Clustering analysis	155
7.4	Spatial and temporal analysis	157
7.5	Identifying Pyromes and Fire Regimes	158
7.6	Conclusions	167
	Bibliography	169
A	Cell2Fire: Cellular-Automata Wildfire Simulator	195

A.1	Cells attributes	195
A.2	Fire propagation dynamics	196
A.3	Detailed inputs	198
A.4	Propagation model validation	199
A.5	Parallel performance analysis	200
B	General Wildfire Spread Model	207
B.1	Additional fire components	207
B.2	Software & Reproducibility	208
C	Downstream Protection Value Detailed Results	209
C.1	Evaluation plots	209
C.2	Experimental Results	211
D	Detailed Results of Extended Framework	216
D.1	Experimental results	216
D.2	Availability of data and material	220
E	Delineating Fire Regimes with AI	221
E.1	Framework and clustering results	221
E.2	Data & Code availability.	239

List of Figures

1.1	A conceptual diagram of the Canadian Forest Fire Behaviour Prediction (FBP) System.	5
1.2	Simulation framework. (1) Raw fuel, topography and fire weather data is pre-processed into Cell2Fire’s format. (2) Cell2Fire calls an independent fire spread model (e.g., FBP), running multiple simulations including proposed treatment or harvesting plans provided by the user to modify the landscape (i.e., cell fuel types). These plans are then evaluated and results are given to provide managerial insights to decision-makers. (3) Finally, outputs are generated and returned to the user.	7
1.3	The elliptical fire spread distribution scheme using the ellipses defined by the Canadian FBP System. At any time t , the backfire will be $BROS \times t$ behind the point of ignition and the head fire will be $HROS \times t$ ahead of the point of ignition of the fire, expanding the ellipse. Then, if the fire spreading cell i reaches the center of an adjacent cell j (e.g., cell $j = i_5$ at time $t(j)$), a new ellipse is generated at time $t' > t(j)$, which triggers a new set of calculations. The wind direction is assumed to be East (i.e., to the right of the main horizontal axis) on the scheme for simplicity of the exposition. In practice, ellipses are oriented in the main wind direction.	8
1.4	Simulation scheme. The send/Receive messages structure facilitates natural parallelization by processing each cell independently. Messages are sent from burning cells (grey squares) to adjacent available cells (white squares). Harvesting periods ($HPeriods$) and Fire periods ($FPeriods$) time scales may differ, with harvesting actions applied every N years while fire activity could be registered in minutes/hours.	9
1.5	Burn probability maps generated after 100 independent wildfire replications in a 40×40 cells landscape located in the Alberta province. Maps generated using random weather scenarios and ignition points (RW-RI, left) and a deterministic ignition (RW-DI, right) illustrate the impact of the stochasticity in the fire growth dynamic. Darker areas indicate higher burn probabilities.	13

1.6	Framework of the inclusion of the different sources of uncertainty in the new system for multiple replications R . The ignition cell j can be specified by the user or generated using a spatial probability distribution. A set of $\omega \in \Omega$ weather scenarios is provided by the user to simulate scenarios for the study area.	14
1.7	FFM problem. Planners perform fuel treatment actions on the landscape over time. They receive the landscape in an initial state (S_0) and execute an action (A_0) following a landscape management policy (π) as a function of the state. This action modifies the landscape in the place where the management is carried out. The objective of this action is to prevent or reduce the spread of fire (F_0), integrated to other managerial objectives like harvesting timber. This action leads to a reward R_1 . In the next period (normally, a year), the landscape begins in a state, given the previous actions and events. This process is repeated until the end of the planning horizon N	16
1.8	Map of the Dogrib instance. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of the area. A hillshade effect has been applied to depict the elevation of the terrain.	18
1.9	Fire growth visualization for sub-instance 2 using Prometheus and Cell2Fire across a heterogeneous landscape with non-flammable cells (mountains, gray cells) and different fuel-types (green and yellow cells).	20
1.10	British Columbia wildfire instances. From the final scars and statistics (left side Prometheus, right side Cell2Fire), we see how the Cell2Fire scars compared to scars produced by Prometheus, reaching $\overline{MSE} = 0.0995$, $\overline{SSIM} = 0.6863$, $\overline{\delta}_{norm} = 16.347$	22
1.11	a) Prometheus fire scar obtained for the region of Dogrib fire area, contrasted with the b) real fire projected into grid format in 2002 and c) Cell2Fire final output.	24
1.12	Map of the Alberta Sub-instance. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of the area. A hillshade effect has been applied to depict the elevation of the terrain.	25
1.13	Pre-treatment analysis. Burn probability maps (left), average ROS heatmaps (center), and consolidated shortest-tree paths (right) are generated for the random weather-random ignition (RW-RI) and random weather-deterministic ignition (RW-DI) experiments. Darker areas indicate higher values (normalized). Fixing the ignition point (RW-DI) is translated into a very focused wildfire dynamic on the eastern side of the landscape, with no fires escaping to the western region of the instance. This results in higher ROS average values in several cells on the east given the dominant land cover in the area and a denser shortest-path tree than the one observed in the random ignition (RW-RI) experiments since all fires are condensed in a unique side of the land.	26

1.14	Post-treatment RW-RI experiments. Management plans (left) are generated using different policies to select which cells to harvest (blue dots). Non-flammable cells are represented by white cells for visualization purposes. Burn probability maps (center) and average ROS heatmaps (right) are generated from 100 independent replications to evaluate and compare different metrics for the selected treatment plans.	27
1.15	Post-treatment BC plans. Management plans (left) for the RW-RI and RW-DI experiments are generated using BC as the main metric to select which cells to harvest (represented with blue dots). Non-flammable cells are represented by white cells for visualization purposes. Burn probability maps (center) and average ROS heatmaps (right) are generated from 100 independent replications to evaluate and compare different metrics for the selected treatment plans.	29
2.1	Surface propagation model elliptical expansion. An example of the elliptical propagation model is shown when the main wind direction aims to the east ($WD = 0^\circ$) and the focus of the ellipse is located at the center of the middle cell. $r(\phi)$ is the estimated rate of spread (ROS) in the ϕ direction with respect to the main wind direction WD	36
2.2	Length-to-breath ratio and ellipses. Examples of ellipses with different LB values are shown, visualizing different potential propagation patterns depending on the shape of the ellipse.	36
2.3	Fitted curves for different LB values (y-axis) as a function of the wind speed WS (x-axis) for the four fire spread models (Kitral, BEHAVE, FARSITE, and FBP). Data points are obtained from each fire spread model, given a set of predefined WS points to use during the curve fitting procedure. The optimal l_1 and l_2 parameters are estimated to maximize the coefficient of determination R^2	39
2.4	Elliptical propagation model components. The diagram illustrates the relationships between the elements of the model and the dimensions of the ellipse including the semi-major axis (a), the semi-minor axis (b), the distance from one focus to the center (c), and the semi-latus rectum (half the length of the chord through one focus, perpendicular to the major axis).	40
2.5	Example of fitted scenario curves for different $HROS$ values (y-axis) as a function of WS (x-axis). We use the a non-linear fitting method to adjust $HROS$ curves to the moisture scenarios. Data points for the fitting procedure are obtained from empirical results introduced in [307].	41
2.6	FireGenerator workflow. The main five steps of the script are depicted.	44
2.7	Samples of Canada, Chile, and Spain original fire scars (left), simulated fire scars post learning (\vec{x}^* - center), and simulated fire scars using default parameters (\vec{x}_0 - right). Individual MSE, SSIM, and δ_{norm} metrics are reported.	53

2.8	Final real and simulated wildfire scars comparison. Historical wildfires for the different instances are shown in even columns and simulated scars after applying the adjustment framework to the general fire spread model are depicted in odd columns.	54
3.1	Framework schematic. (PHASE I) Two modules are integrated: ignition and fire growth (Cell2Fire) models. A cell i is selected for ignition using an ignition model such as random selection, custom probability function, or machine learning models that capture the risk of occurrence in the area. Once a cell is ignited, a weather scenario is selected and a wildfire is simulated. Each fire spreads through the cells following a provided fire-spread model and a shortest-path tree \mathcal{T}_i is obtained. After R replications, a multidigraph \mathcal{GT}_R is generated. Using this information, different risk maps are produced. (PHASE II) This module solves the problem of finding the adjacent cells that minimize expected losses due to future wildfires given specific constraints provided by the decision maker. Finally, the effectiveness of the treatment plans generated is evaluated.	64
3.2	Map of the seven study areas. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of each region. A hillshade effect has been applied to depict the elevation of the terrains.	65
3.3	(a) Planar graph with edges indicating the connection between adjacent cells. (b) Simulated “shortest-path tree” from a heterogeneous instance with arrows indicating the propagation of fire during a replication. (c) Fire shortest traveling times. Red arrows indicate edges where fire was propagated (darker represents longer traveling time in minutes). Axes represent cell ids.	68
3.4	Example of a “shortest-path tree” graph generated from Cell2Fire. Cells are represented by nodes and the edges indicate that fire was propagated from one cell to another during the simulation. Weights $C_{(i,j)}$ include information w.r.t. the traveling times and average rate of spread (ROS).	68
3.5	Shortest-path tree example from Dogrib instance. Arrows are used to indicate the propagation of fire between adjacent cells. Axes represent cell ids. Useful information regarding the ROS [m/min] (left) and travel times [min] (right) are registered at the spreading times between cells. Darker colors indicate larger values – faster rate of spread and longer traveling times, respectively.	69
3.6	Global Propagation Tree obtained from the overlapping of 10 replications (individual shortest-path trees) starting from the same ignition point including uncertainty in the weather scenario $\phi \in \Phi$ in Cell2Fire. Edges frequency usage among replications is indicated with colors, where darker tones represent common propagation patterns. Axes represent cell ids.	71
3.7	DPV comparison for three different “shortest-path tree” graph structures.	73
3.8	DPV versions example. Several combinations of relevant variables of the graph can be used for calculating the node values within the landscape. Its components will depend on the objective and planning horizon of the study.	73

3.9	Multiple treatment plans are tested for different fraction levels (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications e.g., DPV 10, represent the standard deviation of the treatment performance after evaluating the new landscape with 100 independent replications. Adding new simulations from potential wildfires significantly improves the performance of our metric, which is critical to generate effective treatment plans. As an example, focusing on the Arrowhead instance $tf = 0.25$, we observe that using the DPV model, we decrease the average area burned from 281.1 ± 348.0 (DPV 10) to 58.4 ± 139.1 ha. (DPV 100). Here, we notice that even training the DPV model with 10 samples is almost as good as those of BC 100 (second-best), showing the potential of our method. Increasing R , DPV is able to significantly improve its performance, learning faster than any other metric reaching up to 75% of improvement when more data are provided.	78
3.10	“Global Propagation Trees” for $ R \in \{10, 100\}$ replications of the Dogrib RW-RI instance obtained after the overlapping of all individual shortest-path trees in a 100×100 patch (axes). More complex and general wildfire patterns are captured when increasing the number of replications (right) as larger – unexplored with fewer replications – areas of the landscape are covered, allowing the metrics to exploit this information when ranking the cells for the treatment plan.	79
3.11	Treatment plan effectiveness (y-axis) depending on the simulated scenario (ignition point and weather, e.g., DI-DW) and treatment level (x-axis) for Dogrib instance. Fixing the ignition point (DI) to a bounded area (left side) significantly simplifies the problem, obtaining fast convergence toward no losses due to wildfire with all metrics.	80
3.12	Different heatmaps example. Heatmaps obtained after a hundred replications using BP (first row) and DPV (second row) metrics for three different landscape – normalized – value functions (third row). A homogeneous (i.e., all flammable cells with identical value), random (each cell with a random uniform value between 0 and 1), and a volume-oriented value function are used to illustrate the impact on the heatmaps obtained by the metrics, leading to significantly different treatment plans.	81
3.13	Full pipeline example. Heatmaps obtained after a hundred replications using DPV, BC, and BP metrics. Blank cells represent non-fuel sections of the landscape. Highlighted and darker red cells indicate higher metric values (first row). Values are then fed into the protection value problem (PVP), obtaining optimal treatment plans selecting the continuous patch of cells that minimizes the expected wildfire losses (dots connected by edges) for each different metric and treatment fraction level tf such as 5% (second row) and 10% (third row).	82
4.1	Example of layers included visualized in GEE. All layers are obtained for California and consolidated into a 30 by 30 m. multi-band raster.	90

4.2	Framework schematic. Data is retrieved from cloud services and local user inputs. Decision-makers define relevant objectives by analyzing the trade-off between multiple variables. Once the data is processed, multiple simulations are performed to estimate the impact of future wildfires in the landscape. An optimization model is fed with the outputs from the utility mapping and simulation models. Finally, results are evaluated by estimating the average expected losses due to wildfire, as well as analyzing the sensitivity of the optimal treatment plan.	91
4.3	Utility mapper application. (a) The original values of the population density layer (x-axis, popDens) are mapped to the [0,1] interval following an exponential function (y-axis). (b) Density plot for the original variable values. (c) Distribution of the mapped [0,1] feature (called utility) following the applied transformation. . .	94
4.4	Land cover representations. The three case study areas are depicted with a hill shade effect where different colors represent the fuel types characterizing the instances following the Scott & Burgan [307] classification system.	97
4.5	Utility heatmaps for all proposed convex combinations of the main four categories for each instance (columns). The first row represents a balanced combination of all four categories ($\mu_i = 0.25 \forall i$). The suffix <i>dom</i> indicates that the dominant category was weighted by $\mu_{dom} = 0.7$ and the remaining three categories with $\mu_j = 0.1, \forall j \neq dom$	101
4.6	Raw (left) and smoothed (using a Gaussian kernel - right) DPV heatmaps calculated for Napa Valley, Getty center, and Paradise instances using the Access, Forest, and Population density layers as the dominant layers for the <i>NV</i> function, respectively. Lighter cells increasingly highlight the nodes playing a fundamental role in propagating the fire to the rest of the landscape.	103
4.7	(a) Napa Valley instance raw utility (blue) $U(\lambda)$ weighting all categories by identical weights and average discounted utility (orange) $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 25%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 25% of the landscape. Average values are highlighted with red dots.	104
4.8	(a) Paradise instance raw utility (blue) $U(\lambda)$ with carbon as the dominant category and average discounted utility $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ (orange) including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 50%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 50% of the landscape.	105

4.9	(a) Getty center instance raw utility (blue) $U(\lambda)$ with accessibility as the dominant category and average discounted utility (orange) $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 15%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 15% of landscape.	106
5.1	(a) Agents represented as yellow squares can harvest cells (dark green) to contain the fire growth evolution (brown). (b) The true state of the system is represented using a series of matrices and arrays potentially including the current fire progress, topographic and weather information, the position of the agents, and the number of actions required to transform a cell into a fuel break. (c) Agents can move horizontally/vertically or perform an action (stay, harvest the cell to obtain a fuel break).	112
5.2	a) COMA Policy Gradient structure. A central critic computes the advantage function, used as an input for updating the actors' policies. b) Proposed Sub-Groups COMA structure. Thick (red) lines and highlighted text indicate changes in the information passed to the agents and the critic.	116
5.3	(a) Heterogeneous Canadian instances characterized by different flammable land-covers. (b) Specific collaborative tasks are tested in generated instances using non-flammable cells (grey). Given the structure of the landscape, agents require to develop complex collaboration strategies to minimize the expected area burned.	119
5.4	(a) A non-Cautious agent solution is able to keep the fire controlled in the ignition point. (b) The cautious agent solution avoids the eastern side of the land to minimize the risk of getting caught by the fire. (c) Average return comparison when solving the 3×3 instance with a different number of agents N , using a decentralized actor-critic method.	120
5.5	Results samples. (a) Homogeneous instance training comparison ($n = 5$, 1600 cells). (b) COMA and SubG-COMA performance comparison in a heterogeneous landscape ($n = 5$). (c,d) Visualization of SubG-COMA and COMA policies. Agents are highlighted by black circles, harvested cells in dark green, and fire in brown.	121
5.6	Differences between the SubG-COMA and COMA implementations when coordinating agents. Thanks to the local reward function, SubG-COMA agents are able to extract information regarding their contribution to the global reward with less noise, exploiting actions that are useful for the main objective. COMA, on the other side, is able to reach high-quality solutions but agents tend to over-harvest the landscape.	122

5.7	Coordination challenge instances set. Artificially generated instances are used to test the effectiveness of traditional multi-agent DRL algorithms and the proposed SubG-COMA algorithm when the collaboration of the agents plays a crucial role to effectively solve the environment, i.e., minimize the expected losses due to wildfire.	124
6.1	Case study data. (A) Ignition points of wildfires distributed along the national territory of Chile are highlighted. (B) A buffer is generated for each point. (C) A rectangular land-cover image centered at the ignition coordinates is extracted. The total number of wildfires experienced (numbers on the right side of the panel), their main characteristics (e.g., duration, size, and perimeters), and auxiliary variables to assess the predictions of the model (e.g., population and road density) are registered for each rectangular area as part of the dataset metadata.	130
6.2	Summary results obtained for the whole dataset. Representative sample areas classified into the three ignition risk categories, defined from the predicted fire ignition probability thresholds (Section 6.3), are shown in the top-left panel. Performance metrics of the classification algorithm as a function of the predicted fire ignition probability and land-cover proportions for each group (high-, medium-, and low-risk) are shown on the right panels. Resulting attention maps from one hundred randomly selected samples using GradCAM and guided backpropagation (BP) algorithms are shown in the bottom-left panel, highlighting the areas where the model focuses its attention (warmer colors).	136
6.3	Filtered landscapes using different attention levels (full map, top 30%, 50%, and 75% of the densest areas). Attention masks are calculated at different density thresholds to filter the original landscapes, focusing the analysis on the densest sections of the attention maps. Red to blue color represent a gradient from the most to the least relevant zones to classify an image as a positive case. In this example, the network mainly highlights the urban land-cover, represented by areas covered by roads and cities, classifying the image as a fire-positive high-risk landscape.	137
6.4	Risk predictions and characterization using activation maps. Horizontal bars indicate the number of observations at different intervals representing the probability P_i of being classified as a positive class area. Samples are then separated into three categories using P_i thresholds where significant differences are observed in terms of land-cover composition and landscape metrics. Then, these are characterized by calculating zonal statistics within the attention maps hotspots. Zonal statistics are obtained by focusing on high-density areas of the map to characterize and unveil the learning process of the classification model based on the detected spatial/topological structures and interaction between multiple land-covers. This understanding could play a crucial role in defining insightful and practical planning policies incorporating wildfire uncertainty and the mitigation of their future impact.	138

6.5	Zonal statistics at different layers of the network (0, 8, and 19). Statistics are calculated by filtering the original landscapes with the attention maps/masks generated from the GradCAM algorithm, focusing the analysis on relevant regions of the landscape. Warmer areas indicate higher levels of attention to predict the probability of ignition.	142
6.6	End-to-end classification of an example HR landscape. GradCAM, guided GradCAM, and filters visualizations are obtained at the different convolutional layers of the network to improve the interpretability of the model and open the DL black box. The model outputs the probability of ignition associated with each image.	144
7.1	Overall framework of the study. (1) Wildfire data describing individual events in terms of fire-related characteristics such as size, perimeter, duration, and average expansion are collected from products derived from MODIS satellite observations. (2) Data are processed and consolidated into a raster dividing the world into a grid with a resolution of $1^\circ \times 1^\circ$. Annual statistics and features are calculated for each cell, generating numerical (e.g., average fire frequency per time period) and spatial datasets. (3) Statistical methods to analyze multidimensional data are combined with unsupervised learning in order to discover similar groups of cells sharing fire-related characteristics. No explicit spatial components are included. (4) Climatic and socio-economic layers are introduced for each cell in the grid. (5) Spatial density plots are generated for each pyrome, detecting the regions of the world with more observations, assumed to spacially frame a specific regime. Detected fire pyromes and regimes are characterized by climatic and demographic data. An evaluation of the influencing factors is performed for the most relevant areas. A temporal analysis to determine trends and seasonality patterns of fire activity is also carried out. (6) All results and generated datasets are deployed on cloud services and a public-access repository, along with the scripts to reproduce all steps of the study.	152
7.2	Discovering pyromes. Self-organizing maps are useful for summarizing multidimensional fire data and for determining potential groups of similar characteristics. These data are reduced to a two-dimensional grid and samples are organized according to their Euclidean distance. Observations sharing similar characteristics are easily visualized in a topographic map (A) where warmer colors represent widely separated samples and cooler colors depict closely related values. Using image processing algorithms (see Section 7.3), we detect significant potential pyromes/clusters (red circles). The number of observations belonging to each section of the map can be presented in a matrix known as hit-map (B). As an example, we can easily observe the group of cells without fire activity as a large dark blue region (top) and white valley (bottom), representing a significant percentage of the observations.	160

- 7.3 Pyromes' hierarchy. Dendrogram summarizing the hierarchy of the determined fifteen pyromes and similarities between them. As observed, the pyromes can be collapsed into six macro groups sharing fire behavior characteristics, consistent with our statistical results (see Appendix E), where, for example, the pyrome of observations representing extreme and rare events (R10) is clearly independent of other clusters. 161
- 7.4 Fire pyromes. Cells covering the gridded world are classified into the six macro groups determined in the dendrogram (different colors) composed by the determined pyromes based on historical averages of fire characteristics including average annual burning frequency, size, perimeter, duration, daily expansion, and perimeter to area ratio values. The white and light brown backgrounds represent sections of the world where no fire events were registered for the studied period. From the results, certain pyromes cover multiple regions of the globe that do not seem to be related by climatic or demographic conditions. In order to understand their composition, we need to determine and characterize those areas where the observations of each group form a compact and well-defined cluster (via a density analysis) identifying fire regimes associated with specific regions. 162
- 7.5 Regimes and driving factors. Demographic and climatic variables of the most relevant regimes per pyrome are summarized into two independent components using a principal component analysis (PCA) to visualize the regimes' driving factors (A.1, B.1, C.1, D.1, and E.1 panels). A denser population, lower GDP, and harder accessibility characterize regimes located in higher values on the demographic component (y-axis). Regimes located in higher values on the climatic component experience higher average temperatures, higher precipitation levels, and lower PDSI values (x-axis). A shaded region connects regimes within the same pyrome, highlighting the dispersion of the regimes' potential driving factors within each pyrome. The spatial location of the regimes (A.2, B.2, C.2, D.2, and E.2 panels, matching the colors of the regimes in the PCA plots) and the proportion of the dominant land covers per pyrome (A.3, B.3, C.3, D.3, and E.3 panels) are included for reference, providing an overall comprehensive characterization of the most relevant regimes (62 in total) for each pyrome. Detailed results per regime including their dominant land-covers can be found in Appendix E. 164

7.6	Characterizing fire regimes. Five largest spatial regimes are determined and represented with different colors for R1 pyrome after a kernel and contour level analysis. Despite belonging to the same pyrome, dense observation areas are spread in regions with very different climatic and demographic characteristics. In this pyrome, regimes cover parts of the western coast of Canada and Alaska (dark green), the Amazonian forest of Peru and Brazil (blue), the North America great lakes area (green), the eastern extreme of Russia (orange), and the central Asia region (gray). Significant differences can be observed between these regions in (A) location and land-cover distribution; (B) socio-demographic and climatic attributes; and (C-E) seasonal fire characteristics. Land covers observed in the pyrome include savannas (WDS, SAV), grasslands (GRS), conifers (ENC), evergreen broadleaf palmate (EBP), shrublands (OSL), mixed forest (MFS), water bodies (WBS), non-vegetated (NV), croplands (CRO), and permanent wetlands (PWL). Similar comparisons and deeper assessments can be performed for all regimes (see Appendix E).	165
A.1	Dogrib MSE and SSIM hourly evolution (22 hours simulation).	199
A.2	Running times for serial versions. Python’s large instances results are omitted for visualization purposes.	203
A.3	Strong-scaling and speedup factors for OPT version $n = 500,000$	204
A.4	Weak-scaling efficiency (Homogeneous)	205
A.5	Weak-scaling and speedup factors for homogeneous instances, starting with $n = 500,000$	205
C.1	Treatment plan evaluation for treatment fraction $tf \in [5\%, 25\%]$ interval. Different fraction levels are indicated in the (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications.	209
C.2	Treatment plan evaluation for treatment fraction $tf \in [5\%, 100\%]$ interval. Different fraction levels are indicated in the (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications.	210
D.1	Resource allocation sample plans for Napa valley, Getty center, and Paradise (columns) instances for different λ weights (rows) to account for the expected losses due to future wildfires, at a specific treatment fraction tf . Significant variations in the optimal plans can be observed as the λ values are modified to include future wildfire risk into the objective function. Original land cover colors have been modified for better contrast and non-flammable nodes have been removed (white space).	219

E.1	Cluster convergence. Sum of squared distances of samples with respect to the nearest cluster center as a function of the number of clusters k using the K-means algorithm on the data obtained after training our self organizing map. As expected, larger values of k lead to lower SSE values, converging towards 0. We found significant variations in the slope of the function in the [15,20] interval across all tested algorithms.	221
E.2	Dimensionality reduction. Two-dimensional reduction using the t-SNE algorithm with the centroids of the discovered 15 regimes. From the plot, it is possible to observe the clear six macro-groups (highlighted with ovals of multiple colors) and the differences between the regimes.	222
E.3	Pyrome 0. R0 spatial distribution (left) and hot spots (right) representing local fire regimes.	223
E.4	Pyrome 1. R1 spatial distribution (left) and hot spots (right) representing local fire regimes.	223
E.5	Pyrome 2. R2 spatial distribution (left) and hot spots (right) representing local fire regimes.	223
E.6	Pyrome 3. R3 spatial distribution (left) and hot spots (right) representing local fire regimes.	224
E.7	Pyrome 4. R4 spatial distribution (left) and hot spots (right) representing local fire regimes.	224
E.8	Pyrome 5. R5 spatial distribution (left) and hot spots (right) representing local fire regimes.	224
E.9	Pyrome 6. R6 spatial distribution (left) and hot spots (right) representing local fire regimes.	225
E.10	Pyrome 7. R7 spatial distribution (left) and hot spots (right) representing local fire regimes.	225
E.11	Pyrome 8. R8 spatial distribution (left) and hot spots (right) representing local fire regimes.	225
E.12	Pyrome 9. R9 spatial distribution (left) and hot spots (right) representing local fire regimes.	226
E.13	Pyrome 10. R10 spatial distribution (left) and hot spots (right) representing local fire regimes.	226
E.14	Pyrome 11. R11 spatial distribution (left) and hot spots (right) representing local fire regimes.	226
E.15	Pyrome 12. R12 spatial distribution (left) and hot spots (right) representing local fire regimes.	227
E.16	Pyrome 13. R13 spatial distribution (left) and hot spots (right) representing local fire regimes.	227
E.17	Pyrome 14. R14 spatial distribution (left) and hot spots (right) representing local fire regimes.	227

List of Tables

1.1	Accuracy measured by the complement of the mean squared error (MSE) and structural similarity index measure (SSIM) values per hour (6 hours evolution) for the sub-instances. Lower values in SSIM highlight the differences between the cellular-automata and wave propagation approaches as well as the definition of a burned cell between Cell2Fire and Prometheus.	21
1.2	British Columbia simulations summary statistics obtained by comparing the simulated final fire perimeters from Prometheus and Cell2Fire. Mean squared error (MSE) and structural similarity index (SSIM) are reported across all instances. We observe low $\delta_{norm} = X - Y $ where X and Y are the binary BurnGrids matrices produced by both simulators, indicating the high similarity between the results of both simulation approaches.	23
2.1	Elliptical length-to-breadth fitted shape parameters l_1 and l_2 for the proposed general LB expression. Optimal values are obtained for FARSITE, FBP, Kitral, and BehavePlus systems using a non-linear fitting procedure. The coefficient of determination (R^2) is used as the main performance metric for the fitting.	38
2.2	Example of $HROS$ fitted parameters for all moisture scenarios, considering a land cover dominated by shrubs.	42
2.3	Summary of the instances considered including total area in hectares, topographic characteristics (mean and range of the elevation in meters), fire duration in hours, dominant land-covers, and the total number of fuel types per instance.	50
2.4	Best results obtained for each instance. Optimal values for MSE , $SSIM$, and δ_{norm} are indicated (*) and compared with their initial values (0) using \vec{x}_0 as starting point.	55
2.5	Summary of Canadian fuel types adjustment.	58
2.6	Summary of Spanish fuel types adjustment.	59
3.1	Basic description of the instances considered in this study including the total area in hectares, topographic characteristics, and dominant fuel behavior model type and codification following the Canadian FBP System for Dogrib, Arrowhead (AH), Revelstoke (RT), Mica Creek (MC), Glacier National park (GNP), Central Kootenay (CK), and Neptune Peak (NP).	64
3.2	DPV comparison example.	72

4.1	Summary of instances' main characteristics. For each instance, we provide the total area in hectares, the average elevation and its range in meters, the dominant flammable fuel of the terrain following the fuel type layer characterization, the total number of different fuels available in the region, and the number of edges conforming the network used for the optimization model connecting the flammable cells (in any direction).	97
4.2	The expected area burned and expected losses for all utility functions as a percentage of the total instance area and the total utility available (heatmaps) per instance, respectively. Expected values are calculated from $R = 100$ independent wildfire replications, weighting all simulations equally, and without any intervention of the landscape.	102
6.1	Zonal statistics. In this table, we present the statistics for high-risk (HR), medium-risk (MR), and low-risk (LR) images filtered with attention maps at different density levels. We include the total area of the landscape covered by the filter (% area heatmap), the average number of connected components (Ncomp), the sum of all values in the filtered landscape divided by Ncomp (MN), and the Simpson's diversity index (Section 6.3) to illustrate our methodology.	139
6.2	Zonal statistics obtained at different layers of the model. The table summarizes the proportions of croplands (CR), native forest (NF), forest plantations (FP), grasslands (GR), shrublands (SL), and non-vegetated (NV) covers obtained from ten random samples of each three risk categories, filtered using the full attention maps and compared with the proportions of the original landscape. Significant variations across the layers illustrate how the network focuses its attention on specific areas of the image during the classification procedure.	141
A.1	Extract of an hourly weather stream file. Average precipitation (APCP), temperature (TMP), relative humidity (RH), wind speed (WS) and wind direction (WD).	198
A.2	Dogrib accuracy and structural similarity index measure values per hour (22 hours evolution).	199
A.3	Average percentage of the running time parallelizable for different instances. Upper and lower tails are included for completeness. Results obtained by simulating 20 different forests (fuel types, spatial distribution, etc.) for each size n with the same weather conditions.	202
A.4	Speedup factors for small instances for different numbers of threads.	203
A.5	Average speedups for large instances: heterogeneous and homogeneous forests	204

C.1	Average number of burned cells including treatment levels $tf \in [0.05, 0.25]$ for the Alberta region instances. The proposed <i>DPV</i> model is consistently outperforming alternative metrics when trained with $ R = 100$ replications. *Performance difference $\Delta\%$ is calculated for each instance f with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{f_i} - Max(Mean_f))/Max(Mean_f)$ where i is the corresponding row for the f instance.	211
C.2	Average number of burned cells including treatment levels $tf \in [0.05, 0.25]$ for the British Columbia instances. The proposed <i>DPV</i> model is consistently outperforming alternative metrics when trained with $ R = 100$ replications. *Performance difference $\Delta\%$ is calculated for each instance f with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{f_i} - Max(Mean_f))/Max(Mean_f)$ where i is the corresponding row for the f instance.	212
C.3	Learning capacity experiment results for Sub20, Sub40, and Sub100 instances. *Performance difference $\Delta\%$ is calculated for each treatment level tf with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{tf_i} - Max(Mean_{tf}))/Max(Mean_{tf})$ where i is the corresponding row for the tf level.	213
C.4	Learning capacity experiment results for Dogrib, Mica Creek, and Central Kootenay. *Performance difference $\Delta\%$ is calculated for each treatment level tf with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{tf_i} - Max(Mean_{tf}))/Max(Mean_{tf})$ where i is the corresponding row for the tf level.	214
C.5	Learning capacity experiment results for Revelstoke, Arrowhead, and Glacier National Park. *Performance difference $\Delta\%$ is calculated for each treatment level tf with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{tf_i} - Max(Mean_{tf}))/Max(Mean_{tf})$ where i is the corresponding row for the tf level.	215
D.1	Average discounted utility results for Napa instance evaluated from 100 simulations. Results for all λ combinations between the <i>DPV</i> heatmap and <i>NV</i> layers are presented by dominating utility category (column 1) and treatment fraction level (column 2).	216
D.2	Average discounted utility results for the Paradise instance evaluated from 100 simulations. Results for all λ value combinations between the <i>DPV</i> heatmap and the <i>NV</i> layer are presented by dominating utility category (column 1) and treatment fraction level (column 2).	217
D.3	Average discounted utility results for the Getty center instance evaluated from 100 simulations. Results for all λ value combinations between the <i>DPV</i> heatmap and the <i>NV</i> layer are presented by dominating utility category (column 1) and treatment fraction level (column 2).	218

E.1	Pyromes and fire regimes details. Tables E.1-E.15 provide a comprehensive description of all pyromes and regimes. Pyromes are characterized using the inter-annual averages of fire behavior features including frequency, the number of fires, wildfire size [km^2], perimeter [km], duration [days], daily expansion [km^2/day], perimeter-per-area ratio; and the total number of cells and fires. Regimes within the 30% densest areas are characterized by their (1) spatial location; (2) average climatic conditions considering Palmer drought severity index (PDSI), water deficit [mm], temperature [K°], max temperature [K°], and total precipitation [m]; and (3) average socio-economic descriptors including the gross domestic product (GDP) [US\$], population density (population per [km^2]), accessibility (travel time [min] to the nearest densely-populated areas with 1,500 or more inhabitants per km^2), and land-use configuration. Land use includes the following categories: closed shrublands (CSL), croplands (CRO), deciduous broadleaf forests (DBF), evergreen broadleaf palmate (EBP), evergreen needleleaf conifer (ENC), grasslands (GRS), mixed forest (MFS), non-vegetated (NV), open shrublands (OSL), permanent wetlands (PWL), savannas (SAV), water bodies (WBS), and woody savannas (WDS).	228
E.2	Pyrome 1. R1 pyrome and regimes description.	229
E.3	Pyrome 2. R2 pyrome and regimes description.	229
E.4	Pyrome 3. R3 pyrome and regimes description.	230
E.5	Pyrome 4. R4 pyrome and regimes description.	231
E.6	Pyrome 5. R5 pyrome and regimes description.	232
E.7	Pyrome 6. R6 pyrome and regimes description.	233
E.8	Pyrome 7. R7 pyrome and regimes description.	233
E.9	Pyrome 8. R8 pyrome and regimes description.	234
E.10	Pyrome 9. R9 pyrome and regimes description.	235
E.11	Pyrome 10. R10 pyrome and regimes description.	235
E.12	Pyrome 11. R11 pyrome and regimes description.	236
E.13	Pyrome 12. R12 pyrome and regimes description.	236
E.14	Pyrome 13. R13 pyrome and regimes description.	237
E.15	Pyrome 14. R14 pyrome and regimes description.	238

Acknowledgments

I would like to say a big thank you and express my deepest appreciation to my advisor, Professor Zuo-Jun Max Shen. He has fully supported and guided me all these years, trusting and believing in me in every step of this joint journey. I will never forget our fun and long conversations about a myriad of topics, our motivating research discussions, and all the interesting projects we explored together in this dissertation. The first time I entered his office I saw a Professor. Now, I see an amazing friend.

I would like to thank my dissertation committee members Professor Marta C. Gonzalez, Professor Paul Grigas, and Professor Zeyu Zheng for their great support and invaluable suggestions since my qualifying exam. Working with them has been a pleasure. In addition, many thanks to all the Professors of the IEOR department for the great learning experience I had through the program. Same great thanks to the IEOR staff, for their daily support and continuous help.

My deepest thanks to all my co-authors and collaborators: Andres Weintraub, David Woodruff, David Martel, Jaime Carrasco, Alejandro Miranda, Pelagie Elimbi, Marta C. Gonzalez, Jordi Garcia-Gonzalo, and Jose Ramon Gonzalez. Working with such amazing researchers, and most importantly, great persons and friends, has been an incredible and gratifying experience.

There are not enough words to thank all my friends distributed around the world. Many, many thanks to my friends in Berkeley who have made it an enjoyable and unforgettable journey: Pelagie, Alfonso, Riv, Rebi, Bhargav, Sangwoo, Titouan, Ying, Han, Anne, Peter, Jane, as well as to all my fellow PhD students. The “Chilean team”, welcoming me on every trip like if I had never left: Cannobbi, JP, Negro, Maikel, Reid, Andres, Maritza, and Jaime. Last but not least, my awesome friends Jordi and Joserra that has always received me as another member of their families.

Finally, my family. Caro, for starting a new life together and supporting me every second; My amazing parents Sergio & Maricarmen and aunt Mari (Gordis), who have always been there for me; My rock-star sister and brother in law Bel & Yayo, where I will always have my second home; My amazing god parents Ignacio & Mariel and my extended family with Flaca (*mi Chofer!*, where everything started), Coca, Ale, Igna, and Alvarito; My aunt and uncle Carmen & Juanito; My amazing aunt Pili; My super friend Maria Luisa; My brothers Cannobbi, JP, Negro, Maikel, and their families; and my beloved family in France. Without all of you, I would not be here smiling.

Chapter 1

Modeling and simulating wildfires

1.1 Introduction

The effects of global warming on temperature, precipitation levels, soil moisture, and other forest and wildland fire regime drivers have increased, and are expected to continue to increase both the number of and area burned by wildfires around the globe [359]. Wildfires have burned large areas and important infrastructure, thousands of homes and forest resources have been destroyed, and many lives have been lost in recent years. Recent examples include catastrophic incidents in the United States, Canada, Chile, Portugal, and southwestern Australia in the last years [56, 189, 272]. Wildfire occurrences have also resulted in increases in expenditures by forest and wildland fire management agencies (see e.g., [320, 340]). Despite concerted efforts, wildfire growth remains a complex and difficult to model process.

Two of the most important characteristics of a wildfire are its rate of spread (ROS) and intensity, which are influenced by fuel type, fuel moisture, wind velocity, and slope. The Canadian Forest Fire Behavior Prediction (FBP) System includes empirical fire spread rate models that can be used to predict the ROS and the intensity of wildfires based on weather, fuel moisture, time of year, and topographical variables for specified fuel types; e.g., for individual grid cells that contain homogeneous fuel types [125]. However, the FBP System alone cannot be used to predict how a fire will grow across a heterogeneous landscape/grid over time. Spatial fire growth models like Prometheus, a deterministic fire growth simulator, are designed to use FBP spread rate functions to do so [338]. Prometheus is a vector-based fire growth model that is based on an adaptation of Huygens' principle of wave propagation, i.e., the propagation of the fire front is modeled similar to a wave, shifting and moving forward continuously in time and space. It uses spatially explicit fire environment input data concerning topography (slope, aspect, and elevation) and FBP fuel types along with a weather stream and corresponding fire danger rating codes and indices to model wildfire growth [345]. FARSITE is another widely used fire growth simulator [113]. It is based on the U.S. Forest Service's BEHAVE fire behavior prediction system and it is also a vector-based

Huygens' type model. A review of twenty-three simulators that can be used to predict forest fire growth can be found in [255] and a review of new modeling techniques are analyzed in [37]. The two models found to best model the growth of historical fires were FARSITE in the United States and Prometheus in Canada.

Recent years have witnessed a growing interest in the development of detailed cell-based deterministic/stochastic fire simulators. Some of the modeling assumptions that have been adopted include, for example, the use of memoryless probability distributions (Markovian processes) to model the fire spread dynamics [62], homogeneous forests (cells' characteristics are identical), reductions in the number of adjacent cells to which a fire can spread (e.g., from 8 to 6 or 4), and no spotting.

In [2], the authors demonstrated the importance of integrating fuel management with strategic forest management planning models to develop and evaluate FireSmart forest management plans. Such planning calls for integrated forest and fuel/landscape management planning models linked with fire growth models that have well-structured interfaces to facilitate the exchange of data between them. Such models support the iterative re-planning required when strategic plans are modified in response to fires that might occur and the areas that may burn over long planning horizons. In contrast to the previously mentioned tools like Prometheus and FARSITE that are not easily linked to decision-making models chung2015optimizing,rytwinski2010simulation, our simulation model, Cell2Fire, is designed specifically for such use. It is designed primarily to be embedded in a landscape management framework to evaluate fuel-management strategies. Moreover, it can be easily linked with custom optimization models and heuristics to develop good landscape-level fuel management strategies over long planning horizons.

The primary objective of our research and the primary contribution of this project is to develop an efficient and realistic fire growth simulator (Cell2Fire¹) that seamlessly integrates planning models providing a framework that integrates simulation and data-driven decision-making model in the context of landscape management. This is implemented as a high computational performance, open-source tool that exploits parallelism when simulating large-scale fire instances to provide valuable insight to inform both fire and forest management. Ultimately, the research goal is to find ways to inform and support landscape managers in FireSmart forest management activities, allowing decision-makers to efficiently compare, evaluate, and improve their management plans by discovering relevant managerial insights, as we will cover in Chapters 3 and 4. In addition, we aim to provide a tractable methodology that can be used to generate realistic fire scar scenarios to support research concerning fuel management and harvest planning models and algorithms. We use both historical and realistic hypothetical fire instances to validate our simulator and assess its computational performance. We illustrate its use as a decision support tool by applying and evaluating the impact of multiple management plans (e.g., the impact of harvesting a region of the landscape in future wildfire events) on a real landscape.

This Chapter is organized as follows. In Section 2 we introduce the material and methods

¹<https://github.com/cell2fire/Cell2Fire>

of our research and describe the fire growth simulation background, the simulation model, and its integration with our wildfire management decision-making framework. A review of state-of-the-art simulators and details of the computational implementation of Cell2Fire are included. Section 3 presents the results of a case study based on a real forest landscape in Canada and several test instances that were used to validate the simulator output, compare its computational performance with a state-of-the-art simulator, and to illustrate a step-by-step analysis of the integration of the simulation model in landscape management decisions. Section 4 contains our conclusions and thoughts concerning future research needs. Appendix A provides details concerning the case study instances, the computational implementation, and the fire growth model.

1.2 Background

The two methods that have most often been used to simulate fire growth across heterogeneous landscapes are the wave propagation and the cellular-automata approaches. We therefore begin by providing a brief overview of those two approaches.

Wave-propagation model: Huygens

Huygens considered every point on a wave front of light as a source of individual wavelets and described the new wave front as the surface tangential to the circumferences of the secondary waves. The use of Huygens' Principle to simulate fire growth is based on the assumption that the shape of a fire can be represented by a polygon, a plane figure composed of a sequence of straight-line segments forming a closed path whose vertices are a tangential envelope of the elliptical "firelets". Huygens' principle was first applied to the model of fire growth by [298]. Anderson (1982) [25] later developed a simple elliptical model based on Huygens' principle of wave propagation to simulate the growth of grass fires. Richards (1990) [280] then extended this model by deriving a set of partial differential equations to model the growth of fires across a heterogeneous landscape.

Both FARSITE and Prometheus use Richards partial differential equations to propagate each vertex on a fire's perimeter [112, 338]. However, these models differ with respect to the fire danger rating systems and the fuel models used to model fire spread rates. FARSITE uses the fire behavior prediction fuel models developed by [286] and extended by [26] and [306], whereas Prometheus uses the Canadian Forest Fire Danger Rating System and Fire Behavior Prediction System [125].

Cell-based fire growth models

Cellular-automata models that employ a raster-grid of square or hexagonal cells are widely used to model wildfire growth. Fuel and terrain conditions are usually assumed to be homogeneous within each cell in order to simplify basic fire spread rate calculations. The fire

propagates through the grid-cells, typically from a cell's center to the center of an adjacent cell. Each ignited cell behaves as an ignition source that is independent of any adjacent burning cells. To spread the fire from one cell to another, a search mechanism based on an adjacency or spread template is required.

Authors in [186] developed the first computer simulation model to spatially simulate the growth of a small fire. Their model was based on a heterogeneous and discontinuous fuel-type grid but did not account for the effects of terrain and wind. Their deterministic model predicted how long it would take a fire to burn through one square area or cell within a fuel grid when the location of the fire, the starting time, and the grid resolution were known. Travel times were calculated using fixed rates of spread (based on the fuel type and the spread index for the day) and fixed spread directions from the burning cell. Later, O'Regan et. al (1973) [246] developed a method for using directional rates of spread to predict fire growth. They also rewrote the original model for use on what was then a large computer, to simulate the growth fires of up to 15,000 ha in size.

In [245], the authors developed a model for average directional rates of spread and [187] modified that model to accommodate variation in hourly wind conditions. Todd (1999) [334] adapted the model presented in [187] to create an eight-point symmetric fire growth model called Wildfire, which incorporates FBP System spread rates. The features and functionalities of the wildfire model were assessed and considered during the design of the Prometheus model [338]. The authors in [62] developed a stochastic model of fire spread using a lattice Markov chain model in which they associated probabilistic transition functions with each cell. Each of these cells interacts with its four nearest neighbors and a cell transitions from unburned to burning depending on the state of its neighboring cells. The use of a simplified cellular-automata model describing the dynamics of fire spread on a heterogeneous landscape accounting for weather factors (wind speed and direction) as well as the type and density of vegetation was used to successfully model the Spetses Island fire (Greece) after tuning the main parameters of the simulator [15].

Recently, [28] released a fire growth simulator designed to assist civil protection and fire management agencies with a case study in the island of Sardinia, Italy. Their model uses the level set technique (see [135]) and the [286] fire behavior model. Such models are designed to be used in near real-time to inform the on-going management of specific fires that are subject to active suppression action.

One alternative to the cell-based approach calls for probabilistic spatially explicit fire scenarios in the form of burn scar maps that describe what portions of the forest might burn in the future and the probability that each of those scenarios might be realized using, for example, methodologies like the one described by [193]. Such scenarios could be provided to planners engaged in traditional scenario planning exercises (e.g., [228]) and incorporated in stochastic programming models (e.g., [171]) designed to support such planning.

Fire behavior models: The Canadian Fuel Behavior Prediction System (FBP)

The Canadian FBP System is a set of empirical models that can be used to predict fire spread rate, fuel consumption, and fire intensity within homogeneous spatial units (i.e., cells) as functions of fuel type, slope, fuel moisture, and current weather expressed in terms of the Canadian Forest Fire Weather Index System (FWI) codes and indices [125] (Figure 1.1). It includes fuel models that are used to classify vegetation into 17 fuel types that collectively represent most of the major forest cover types in Canada. In the context of landscape management, outputs generated from the FBP system can be used to inform the development of landscape management plans. This allows the planner to incorporate fire behavior outputs like the average ROS, expected flame length, and fire intensity, among other features, to evaluate the impact of proposed plans to mitigate potential effects of future wildfires in the area of interest.

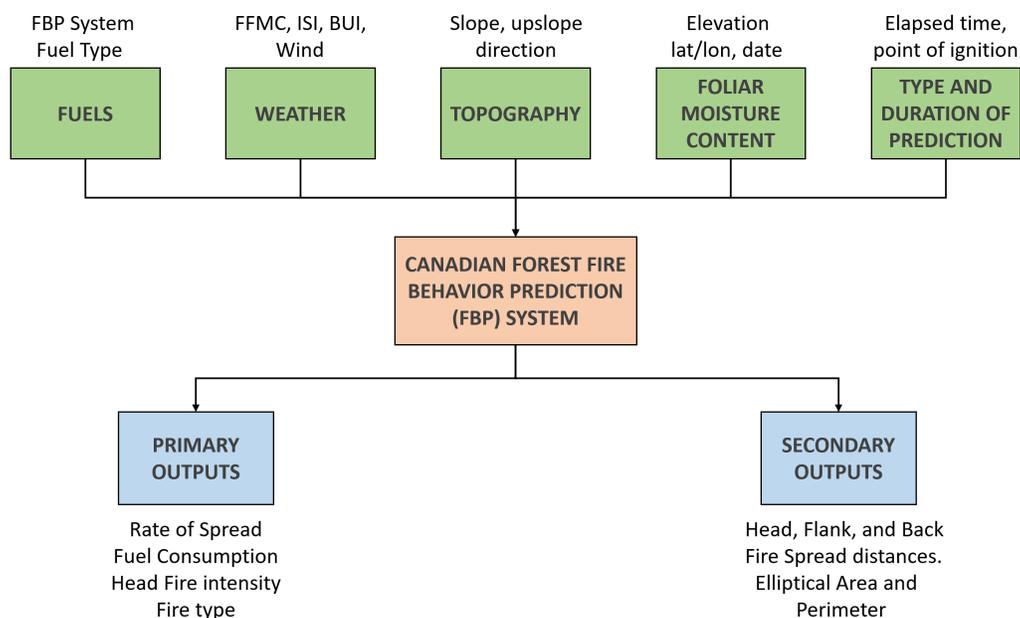


Figure 1.1: A conceptual diagram of the Canadian Forest Fire Behaviour Prediction (FBP) System.

1.3 Cell2Fire growth simulator

Cell2Fire is an open-source cell-based fire growth simulator developed using Python and C++ for laptop or desktop computers as well as on High-Performance Computer systems.

It allows a user to simulate fire growth across a grid that represents a real forest landscape using fire environment variables such as the fuel type, topographic features of each cell, fire ignition points, and weather scenarios (Figure 1.2).

A forest landscape is divided into a rectangular region comprised of rows and columns partitioned into a series of square cells, all of which are the same size. This generates a grid in which the cell size depends on the desired spatial resolution and the granularity of available data. Each cell represents a specific portion of the landscape and has two information layers that pertain to its topographic and fuel characteristics. Those layers define the characteristics of each cell, allowing the simulator to treat them as individual objects that can interact to model fire growth in highly-scalable parallel implementation. Algorithmically, Cell2Fire simulates the growth of fire by tracking the state of all cells as the model progresses through discrete equally-spaced time steps (e.g., seconds, minutes). The status of the fire and all the cells on the landscape are updated (see Section 1.5 and Appendix A for more details) at the end of every time step. Once an ignition point has been specified, the fire is ignited. During each time step, the fire may spread along the axes emanating from the center of each burning cell to its neighboring cells. The predicted FBP system Head Rate of Spread (HROS), Flank Rate of Spread (FROS), and Back Rate of Spread (BROS) are used to model elliptical fire growth within each cell with the focus of the ellipse at the center of the cell (see details in Appendix A). The geometry of the ellipse for every burning cell is computed and used to predict the fire spread rates along the axes emanating from the center of each cell (Figure 1.3).

A signal/message is transmitted and received instantaneously by any adjacent cells whose center is reached by the fire. In the present implementation, it is assumed that each cell has at most 8 adjacent cells (rectangular grid). These are the only neighbors considered because the simulation time step is set small enough to ensure that the fire cannot spread beyond adjacent cells in one time step. Shorter simulation time steps result in longer running times but more accurate spatial (and thus, with incidence in temporal evolution) simulations. We note that there is also a simplifying assumption because when fire enters a cell from one of its neighbors, that neighbor's cell characteristics are used to model fire spread within the destination cell until the fire reaches its center. At that point, the characteristics of the destination cell are used to model further fire spread within the destination. This approximation results in computational efficiencies because a cell does not have to recognize multiple fuel types.

When the fire reaches the center of a cell, it receives a message. It then initiates the calculation of the ROS towards its eight neighboring cells based on its characteristics and the current weather (i.e., fire danger rating indices, wind velocity, and slope). If the calculated ROS (for each of the eight axes) is greater than a user-specified threshold parameter (representing the minimum ROS at which the fire effectively propagates within and between cells), its state is updated. In our experiments, that threshold was set to zero; however, the software provides the ability for fire behavior experts to set a non-zero threshold for particular simulations. This could be used, e.g., to simulate wildfires where suppression actions were applied, in an attempt to represent their effect in the fire growth model. If the ROS exceeds the threshold, the main ROS values are calculated by the FBP System module and

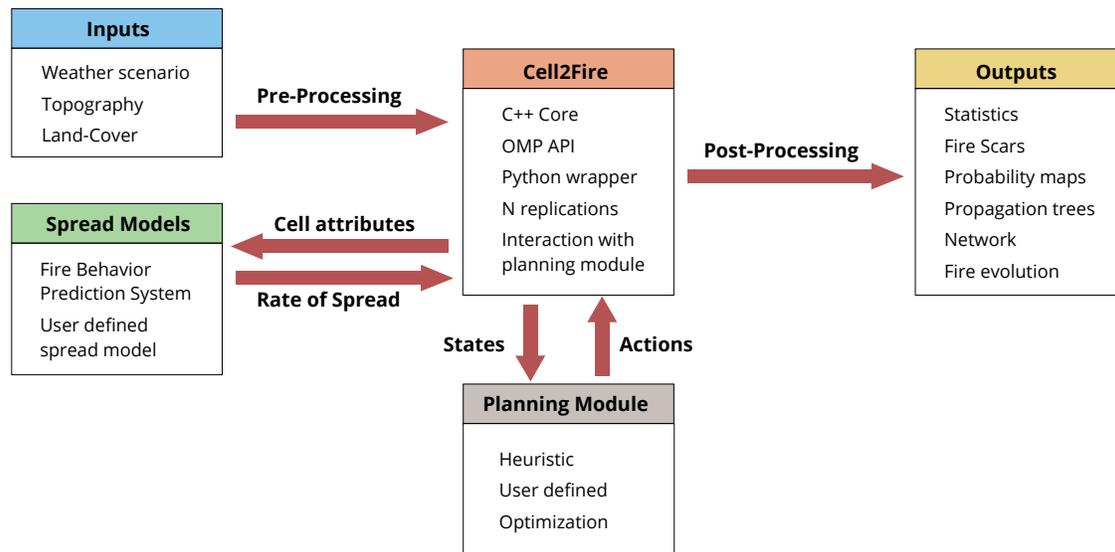


Figure 1.2: Simulation framework. (1) Raw fuel, topography and fire weather data is pre-processed into Cell2Fire’s format. (2) Cell2Fire calls an independent fire spread model (e.g., FBP), running multiple simulations including proposed treatment or harvesting plans provided by the user to modify the landscape (i.e., cell fuel types). These plans are then evaluated and results are given to provide managerial insights to decision-makers. (3) Finally, outputs are generated and returned to the user.

fire progress is predicted along the available axes.

Then, the fire’s progress is updated at the end of each fire time period by examining the state of all active burning cells. Once no adjacent cells are available or a burn-out criterion (See assumption (A5) in Section 1.3) has been satisfied, the cell becomes inactive and is excluded from further simulation steps. This process is repeated until a fire-ending event state is reached, i.e., the total number of fire burning periods (e.g., hours) have passed or there are no more cells available to burn.

The full simulation is presented in Figure 1.4 which illustrates a forest with 9 cells where no harvesting has taken place ($HPeriod_1$). A fire is ignited in cell 4 and it can spread to other cells during the following fire spread periods. If no messages are sent to neighboring cells based on the current environmental conditions (burned-out conditions) or the maximum simulation time for the current fire has been exceeded, the simulation advances to the time the next fire is ignited on the landscape (randomly generated or user-provided) or stop the simulation. A general pseudo-code of the simulation steps is provided in Supplementary Material, Algorithm S1.

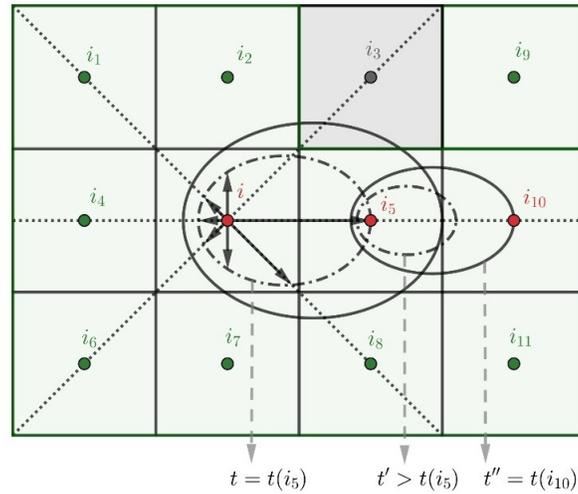


Figure 1.3: The elliptical fire spread distribution scheme using the ellipses defined by the Canadian FBP System. At any time t , the backfire will be $BROS \times t$ behind the point of ignition and the head fire will be $HROS \times t$ ahead of the point of ignition of the fire, expanding the ellipse. Then, if the fire spreading cell i reaches the center of an adjacent cell j (e.g., cell $j = i_5$ at time $t(j)$), a new ellipse is generated at time $t' > t(j)$, which triggers a new set of calculations. The wind direction is assumed to be East (i.e., to the right of the main horizontal axis) on the scheme for simplicity of the exposition. In practice, ellipses are oriented in the main wind direction.

Modeling assumptions

We can summarize the main simplifying assumptions upon which Cell2Fire is based as follows:

- A1) The growth of the fire depends on the ROS from burning cells towards their neighboring cells. We assume that a cell is ignited when the fire reaches its center and conditions for burning are met (see A5). Each cell has at most eight neighbors.
- A2) The ROS along the eight principal orthogonal axes of each burning cell are calculated using the Canadian FBP System as functions of the weather, slope, and fuel characteristics of each cell. The major axis of each ellipse is aligned in the HROS direction and the BROS is the opposite direction. The FROS is perpendicular to the HROS and BROS axis. We note that alternative fire spread models could be used in lieu of the Canadian FBP system, as discussed in our Conclusions section.
- A3) Each cell that burns serves as a new source of fire. Fire spread occurs at the cellular level and cell size depends on the spatial resolution and corresponding availability of

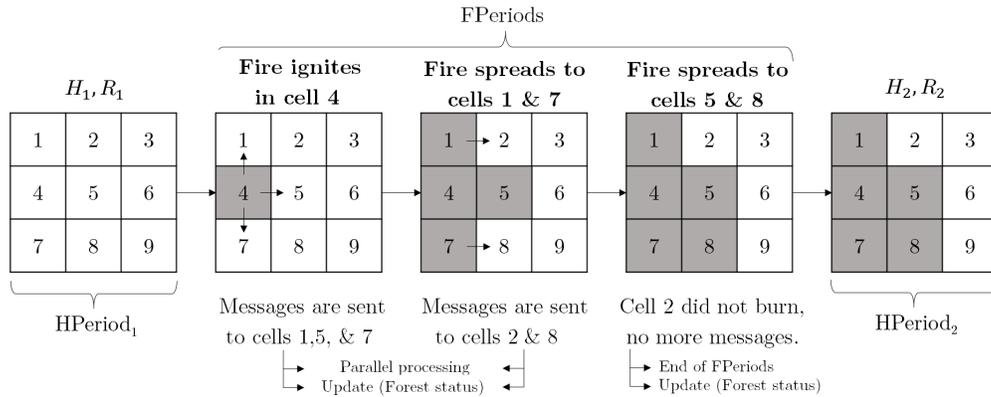


Figure 1.4: Simulation scheme. The send/Receive messages structure facilitates natural parallelization by processing each cell independently. Messages are sent from burning cells (grey squares) to adjacent available cells (white squares). Harvesting periods (*HPeriods*) and Fire periods (*FPeriods*) time scales may differ, with harvesting actions applied every N years while fire activity could be registered in minutes/hours.

fuel and topography data.

- A4) The effect of fire suppression action is not modeled as it is beyond the scope of this research.
- A5) There are two sets of conditions for modeling the termination of fire growth in Cell2Fire: the cellular level and general fire evolution conditions. At a cellular level, each cell becomes unavailable (i.e., burned and can no longer serve as a source of fire) if (i) the ROS along the HROS axis is less than some empirical threshold $\delta > 0$; (ii) the cell does not have any adjacent cells that are available to burn; (iii) the residual fuel available in a cell is not sufficient to support fire spread (implicit in the FBP system), or (iv) a user-defined head fire intensity (HFI) threshold is provided and the HFI is below that threshold. Regarding the general fire dynamics, the total duration of the simulated wildfire event is determined by (1) the maximum number of hours of burning per day – a season-dependent constant [261], or drawn from a probability distribution – and (2) the total fuel remaining (available cells) in the forest (i.e., when it burns out).

Main Inputs and Outputs

As is the case with other state-of-the-art fire growth simulators, the Cell2Fire model requires a number of inputs including a minimum set of data layers that define an instance/forest. The relevant inputs needed to simulate the growth of a fire using Cell2Fire are as follows:

- i) Forest raster data: gridded forest attribute files that specify the number of cells in the forest, their geographical coordinates, and information concerning each cell including its fuel type, elevation, slope (% and azimuth), and the degree of curing of grass (optional).
- ii) Fuel type dictionary: Fuel type codes and descriptions that match the Canadian FBP System fuel types. Custom dictionary files including user-defined fuel types matching an alternative fire spread model can be provided.
- iii) Ignition points: An optional file that that specifies the cell(s) in which fires are to be ignited during the simulation, paired with their corresponding ignition time periods.
- iv) Weather stream: Hourly weather records from one or more fire weather stations located near the area of interest. Weather scenarios using the FBP system include the date-time, precipitation, temperature, wind speed/direction, relative humidity, scenario ID, as well as the hourly fire danger rating codes and indices (FFMC, DMC, DC, ISI, BUI, and FWI) of the Canadian Forest Fire Danger Rating System [329] (see Appendix A).

Once a simulation run has been completed, the following outputs are available:

- i) Burn-Grids: Files in which 1s indicate burned cells and 0s indicate those cells that have not burned. That data can be used to compare our fire growth predictions with predictions produced by other fire growth simulators as well as to generate burn probability maps, or to generate confidence intervals for cell-specific burn probability estimates.
- ii) Plots: Initial forest state, fire scar evolution, and message sending/receiving can be visualized by a series of plots generated after the simulation run has been completed.
- iii) Statistics: Final status of the forest with relevant information. This includes shortest propagation paths, fire behavior data by time step (e.g., average ROS values, fire length), expected fire scar perimeter/size/location, and a series of statistics across all tested scenarios (e.g., including multiple ignition points, weather streams).

1.4 Comparison Methodology: Fire growth validation

In this study, we compare predicted burn scars/perimeters by measuring the difference between fire perimeters generated by Prometheus (baseline) and Cell2Fire simulations using three metrics. The simplest and most widely used full-reference quality metric is the mean squared error (MSE), which objectively quantifies the deviation from a known pattern. However, two distorted images with the same MSE may have very different types of errors, some of which are much more visible than others. Thus, as we are interested not only in the difference of the MSE but also in structural information, we use a measure of similarity suggested

in [373] and denoted by SSIM (structural similarity index). Finally, we also include the Frobenius norm of the difference between two perimeters X, Y , $\delta_{norm} = \|X - Y\|$, focusing on the spatial comparison between fire perimeters.

We compare the evolution of Cell2Fire and Prometheus fire perimeters on a period-to-period basis (where a period represents one hour) in order to measure the differences in fire propagation rates. We denote by $PromGrid^t$ a 0-1 Matrix at time t , which represents the fire scar obtained with Prometheus, where $PromGrid_{ij}^t$ is equal to 1 if the cell (i, j) was burned at time t and 0 otherwise. Analogously, we define the fire scar obtained by Cell2Fire as $Cell2Grid^t$. We note that Prometheus fire perimeters are used as references to validate our fire growth simulation model comparing it to an acknowledged state-of-the-art simulation model. However, its outputs do not necessarily match historical fire perimeters due to a series of approximations/limitations (e.g., land cover data availability, not incorporating suppression efforts applied during the wildfire, etc.).

Below, μ_{X^t} , μ_{Y^t} , σ_{X^t} , σ_{Y^t} and $\sigma_{X^t Y^t}$ represent the means, standard deviations, and cross-covariance for fire scars X^t and Y^t respectively, and C_1 , C_2 , and C_3 , are internal parameters of the metric [373]. The methodology is as follows:

1. Choose an ignition point for each instance and run Prometheus for T time periods (e.g., hours). Thus, we obtain $X^t = PromGrid^t$, $t = 1, \dots, T$.
2. Choose the same ignition point as above and run Cell2Fire for T time periods obtaining $Y^t = Cell2Grid^t$, $t = 1, \dots, T$.
3. Calculate, for all t :
 - a) Mean Squared Error:

$$MSE(X^t, Y^t) = \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m |X_{ij}^t - Y_{ij}^t|^2,$$

to measure average of the squares of the pixel differences of the fire scars (with X_{ij}^t the (i, j) component of the matrix X^t).

- b) Structural similarity Measure:

$$SSIM(X^t, Y^t) = \frac{(2\mu_{X^t}\mu_{Y^t} + C_1)(2\sigma_{X^t Y^t} + C_2)}{(\mu_{X^t}^2 + \mu_{Y^t}^2 + C_1)(\mu_{X^t}^2 + \mu_{Y^t}^2 + C_2)},$$

to measure the change in structural information between the fire scars obtained from the two simulators: Cell2Fire and Prometheus.

4. Measures analysis: MSE and SSIM over times t , and δ_{norm} for the final scar.

1.5 Wildfire Management

One way to mitigate the impact of fire and protect our communities and natural resources is through *Forest Fuel Management* (FFM). Activities such as harvesting, prescribed burns, allocation of firebreaks, commercial harvesting, and thinning can reduce the detrimental impact of fire [3, 209, 240, 252]. In addition, these actions could also benefit the diversity of species [312] and improve the health of the forest ecosystem, among many other benefits [116, 157, 230]. Land planners face difficult questions concerning what, where, when, and how to manage fuels. Moreover, there are few computational tools that address this challenging problem [87].

In general, for a given planning horizon, planners must decide where to locate a treatment depending on the state of the forest. The condition of the forest depends on the previous efforts of the land managers (e.g., the actions taken during previous years) and possible modifications of the vegetation due to wildfires. Now, where and when the fires will occur and what will be their severity are stochastic events. Therefore, the question we propose to answer is: What fuel management activities might best minimize expected future losses? However, regardless of who or how we make the decisions, we require an efficient modeling framework and computational tool to test the effectiveness of alternative plans.

In the next subsections, we show how Cell2Fire can be used to support the development and evaluation of FireSmart forest management plans, focusing on fuel treatment decision-making. We first indicate how the system works to perform multiple stochastic simulations on the study area, with the purpose of calculating relevant outputs such as Burn Probability maps. Then, we show how the integrated decision-making model is capable of modifying the state of the treated cells (with some predefined policy) and measure the impact on cells affected by fires in multiple replications, allowing the evaluation of the effectiveness of different policies in a quantitative framework. Following this line, we will discuss in the next Chapters how the system can be used to generate inputs for Stochastic Programming (SP), Reinforcement Learning (RL), or Simulation Optimization (SO).

Burn Probability Maps

Burn probability maps (BP-maps) are commonly used to assess the likelihood of burning [5, 7, 120, 239, 257]. Software such as Burn-P3, FSPro, and FlamMap [120, 257, 260] can calculate these values using fire growth algorithms to produce high-resolution spatial estimates. Although not the most effective technique in the context of fuel treatment to mitigate future fire losses, potentially leading to treatments that rarely intersect with wildfire occurrences [44], its simplicity and easy interpretation become an advantage for a first assessment and to obtain a baseline to compare with more complex approaches such as network-based prioritization metrics and simulation-optimization techniques [2, 240, 252, 291]. The system proposed in this study includes this capacity and in what follows we describe how it works.

The general procedure for calculating a BP-map consists of generating multiple simulations on a landscape, in which a cell/area is selected at random on each replication, and from

that point, a simulated fire is generated following given or randomly selected meteorological conditions. For example, Burn-P3 (Probability, Prediction, and Planning) [261] combines a deterministic fire growth model (Prometheus) and spatial data for forest fuels and topography with probabilistic fire ignitions and spread events derived from historical fire and weather data. The components of the model include the location and frequency of ignitions, the rate at which fires escape the initial attack and become large wildfires, the number of days on which each fire achieves a significant spread rate, the weather conditions associated with these spread event days, and the deterministic fire spread evolution. Other tools with similar features include FlamMap [114] and Fire Spread Probability (FSPro) [119].

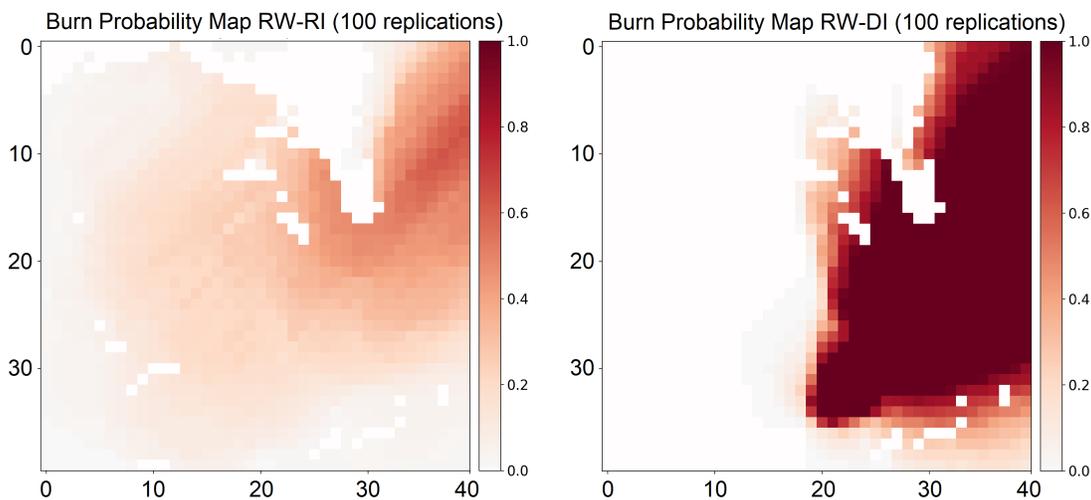


Figure 1.5: Burn probability maps generated after 100 independent wildfire replications in a 40×40 cells landscape located in the Alberta province. Maps generated using random weather scenarios and ignition points (RW-RI, left) and a deterministic ignition (RW-DI, right) illustrate the impact of the stochasticity in the fire growth dynamic. Darker areas indicate higher burn probabilities.

In order to produce the BP-maps with our system shown in Figure 1.5), we use methods similar way to the systems just described. The three different sources of uncertainty included are: (1) Ignition point(s) selected via a user-defined probability distribution or a simple uniform approach for each period of the planning horizon. (2) A coefficient of variation (cv_{ROS}) capturing the stochastic aspects of the ROS predicted by the fire spread model and accounting for its inherent approximation error, allowing the user to obtain different fire scars by including uncertainty in the fire dynamic. (3) A set of user-generated weather stream files (scenarios) with specific probabilities that can be provided to Cell2Fire, obtaining simulations with different weather conditions such as wind direction/speed, precipitation, among others (Figure 1.6).

The general scheme is shown in Figure 1.6. First, we use the ignition locations as a user-defined parameter, giving control over the pattern of ignitions on the landscape. These points can be selected randomly or deterministically before the simulation begins. The latter can model, for example, the lightning strikes over the forest. When lightning strikes a cell i at time t , there is a “ignition probability” (defined by the user) that the fire will start or not. On the other hand, Cell2Fire incorporates a *stochastic weather module* that draws a weather scenario ω from a set provided by the user, Ω . At this point, the user must previously construct the scenarios using, for example, historical data to produce representative and different variations of the weather, where the latter has the same format as those of Prometheus simulator and Burn-P3 (see Appendix A). In addition, users can assign a probability to each scenario in order to generate relevant statistical outputs while providing more flexibility when analyzing the impact of potentially extreme but rare (with low probability) events.

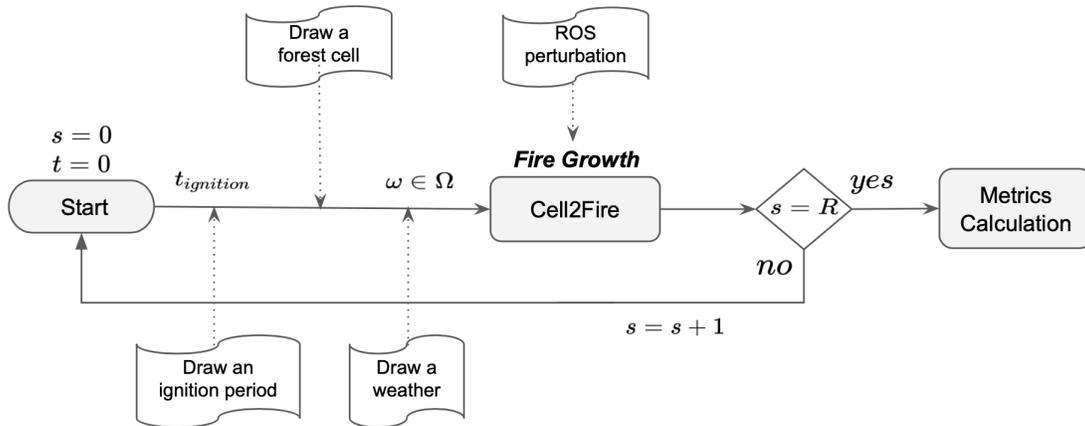


Figure 1.6: Framework of the inclusion of the different sources of uncertainty in the new system for multiple replications R . The ignition cell j can be specified by the user or generated using a spatial probability distribution. A set of $\omega \in \Omega$ weather scenarios is provided by the user to simulate scenarios for the study area.

Although the methodologies differ somewhat, all tools incorporate the main components included in our study: random selection of ignition points, selection of weather scenarios, recording of burned cells, and repeating the process a number of times (Figure 1.6). The main difference with our approach is that they do not consider uncertainty in the ROS, therefore, in the event of randomly selecting the same ignition point and the same weather scenario, their systems produce the same fire scar. In addition, our system can record the cells burned in previous iterations/simulations, since cells have states, unlike Burn-P3 where this information is not recorded as the fire moves on a continuum over the landscape.

Integration with fuel management decision making

The intended users of Cell2Fire are researchers working to develop methods to support decision-making, but in the interest of cleaner exposition we write this section to refer directly to decision-makers. Decision-makers can implement, analyze, and evaluate the impact of different landscape management plans under wildfire uncertainty using our proposed framework. Cell2Fire was designed so that spatial simulation of fires can be easily integrated with data-driven (i.e., with dynamic data representing the state of the vegetation and weather conditions) fuel management decision-making models. It can be used as a pure simulation tool or embedded in a landscape management framework. The latter allows decision-makers to evaluate and observe the potential outcomes and impact of their fuel management strategies or link the tool with optimization models to develop efficient landscape-level fuel management strategies over long planning horizons (to be covered in future extensions). These characteristics can make Cell2Fire a crucial evaluation tool during the planning process.

Let S be the set of states of the cells in the landscape. The state represents the landscape condition that could be modified by actions taken by decision-makers and stochastic wildfire events (Figure 1.7). A cell can be in one of four states: “Available”, “Burned”, “Treated”, or “Non-Fuel”. The label “Available” indicates that the cell contains a flammable fuel type; “Burned” indicates that the fire has consumed the fuel available and passed through the cell (in this first study we assume that 100% of the fuel is burned); and “Non-Fuel” is a non-flammable fuel type such as rivers, lakes, or rocks. The “Treated” state is provided so that the system changes the state of the cell when the planner locates a fuel treatment in it. A cell can be assigned more than one treatment, depending on the characteristics of the type of forest fuel (some cells may not be available for all treatments) and the compatibility between the treatments. We define $S_0 \in S$ as the starting state of the landscape.

A landscape management plan (denoted π , a function of the state S_t) consists of a series of cells to be managed (e.g., harvested) at a specific period of time within the simulation horizon. Depending on the type of managerial action modeled, the original land-cover of the selected cells is modified to represent the treated landscape. For example, management plans involving harvesting a set of cells will transform these flammable areas into non-flammable ones, which can act as firebreaks. Thus, decision-makers will be able to evaluate, for example, which areas to harvest in order to minimize the expected impact of future wildfires, supported by thousands of simulations in practical running times. This, while satisfying operational and budgetary constraints in a flexible but realistic simulation environment. On the other hand, intermediate treatments such as thinning or pruning can be modeled by feeding the simulator with an alternative land-cover layer to represent the treated landscape or an auxiliary table indicating the impact of different actions for each type of fuel cover and the cells to treat/manage. We use $\mathcal{A}(S)$ to represent the set of actions allowed in state $S \in S$. In this first version, we consider only harvest actions or firebreaks. Firebreaks are designed to stop outright or impede a fire’s progress. We will assume that once a cell is treated or affected by fire, it does not recover, that is, it does not re-grow its vegetation at a

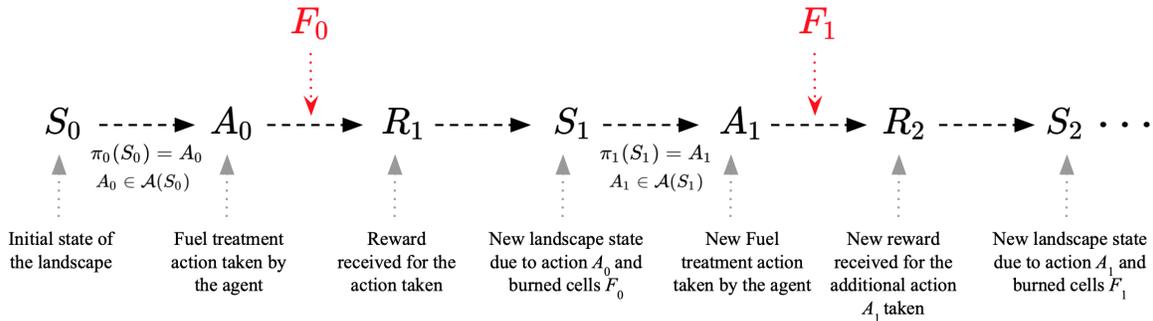


Figure 1.7: FFM problem. Planners perform fuel treatment actions on the landscape over time. They receive the landscape in an initial state (S_0) and execute an action (A_0) following a landscape management policy (π) as a function of the state. This action modifies the landscape in the place where the management is carried out. The objective of this action is to prevent or reduce the spread of fire (F_0), integrated to other managerial objectives like harvesting timber. This action leads to a reward R_1 . In the next period (normally, a year), the landscape begins in a state, given the previous actions and events. This process is repeated until the end of the planning horizon N .

later time-step. This assumption can be relaxed by integrating Cell2Fire with forest growth simulators or by updating the state of the previously harvested cells after specific time-steps (representing their growing period) to certain fuel types, leading to new ROS calculations. This topic can be explored in future studies.

Users have multiple options to model both harvesting and treatment plans, including (1) directly providing the set of cells to treat every simulated period; (2) selecting one of the heuristics provided by our simulation-optimization model; or (3) implementing their own algorithm. Using the first and third approaches, managers can easily evaluate the performance of their management plans under wildfire uncertainty. Given a fixed plan or custom logic, they can simulate multiple wildfires using historical data from the region of interest to evaluate the potential impact of future wildfires in the area, observing the pre and post-treatment results. Alternatively, Cell2Fire incorporates a series of out-of-the-box planning heuristics to provide quick evaluation baselines to the decision makers as well as a guide for users to implement their own logic within the decision-making module. Starting from one of the simplest and most studied methods [257], Cell2Fire can generate automatic treatment plans selecting those cells with highest wildfire susceptibility (burn probability maps, Figure 1.5) based on multiple wildfire replications under the given conditions. It can also select those cells that experience the highest ROS during the simulations, in an attempt to minimize the average ROS within the landscape to decrease the out-of-control wildfire propagation risk. Given an input layer representing the economic value of each cell, e.g., the value of the timber after harvesting a cell, Cell2Fire can generate harvesting plans

that select those cells that maximize the total expected revenue. This revenue model can explicitly incorporate wildfire uncertainty by weighting it with a wildfire behavior metric (e.g., burn probability) or it can completely ignore its potential impact (a myopic approach). Alternatively, managers can provide custom metrics as input layers to generate optimal harvesting plans following this same logic. In this context, $\mathcal{R}(S, A, S')$ denotes the immediate reward earned in state S when action A is selected in it and, as a result, the system transitions to state S' . For example, a reward can quantify the available surface or the volume of harvested timber after taking action A and experiencing the next wildfire F .

In addition, our decision module is integrated with a powerful network flow package [149], allowing us to generate a series of management plans based on network/centrality metrics and combine them with user-provided layers to establish landscape utility functions (heatmaps) to be optimized during the simulations. These metrics, that aim to identify the most relevant nodes within a network, could play a crucial role in the development of effective management plans by detecting those nodes where the fire tends to propagate faster, more frequently, with the higher intensity, or that maximize the propagation of the fire to multiple areas of the network, among many other relevant measurements. In addition, Cell2Fire provides access to node influence metrics aiming to assess the influence of each node in the network and the potential impact of removing/treating them, sharing concepts of previous works such as [2, 335]. Examples of mathematical models to identify and optimize fuel treatment locations, incorporate the role land-owners in fuel treatment activities via a cost-share approach, and maximize the effectiveness of the initial attack in the context of wildfire suppression can be found in [46, 274, 275]. Finally, a random treatment heuristic is included where available cells are selected at random in the landscape, establishing a baseline for comparison purposes. These management plans are flexible, meaning that decision makers can easily modify their main input parameters and add specific constraints to obtain realistic plans, as we will cover in the following Chapters. This includes budgetary or maximum area treated constraints, satisfy adjacency/connectivity constraints (selecting a set of continuous or fragmented set of cells to treat), prioritize specific areas (e.g., protected habitats), as well as custom ones provided by the landscape manager (see [240, 252] for an application of this framework)

1.6 Experimental Instances

We used three sets of fire instances to compare the Cell2Fire simulations with those generated by Prometheus. We wanted avoid introducing comparison bias with respect to real fire scars where the intensity of suppression efforts could be uncertain. The instances are: (1) Dogrib fire landscape in Alberta province, (2) Dogrib fire sub-instances set, and (3) real landscapes in province of British Columbia with hypothetical wildfires. The largest instance from the second set is also used to illustrate the applications of the proposed management module. For the experiments, we used a simulation time step of one minute.

Dogrib fire instance

We chose to model this particular September-October 2001 fire in the province of Alberta due to the vast amount of documentation and observed data available (e.g., weather conditions recorded from The Yaha Tinda Automatic station, and demographic/topographic data collected from the area). It also contained a representative set of different fuel-types documented in the Canadian FBP system (see Figure 1.8). We divided the landscape into 79,611 100×100 [m^2] cells, used the Dogrib fire’s ignition point located at (51.652876° , -115.477908°) and started growing the fire at 13:00 hrs on October 16, 2001 so as to capture the major fire run (ninety percent of the total 10,216 ha of area burned). The original ignition point is translated into an ignition area (cell) in Cell2Fire at its coordinates. This instance is provided with Prometheus².

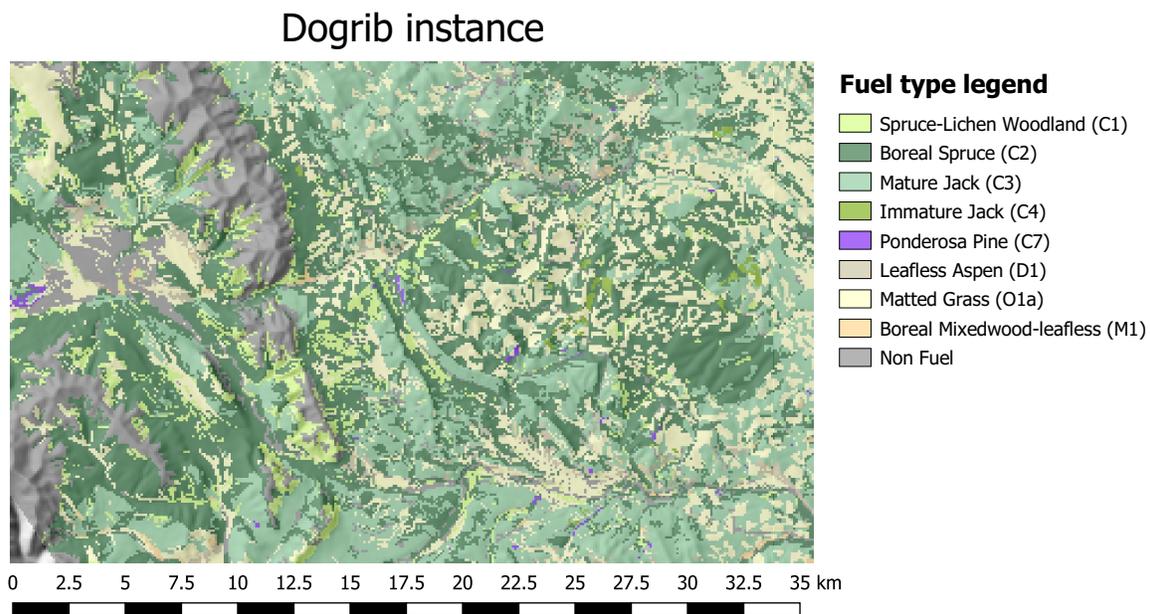


Figure 1.8: Map of the Dogrib instance. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of the area. A hillshade effect has been applied to depict the elevation of the terrain.

Dogrib sub-instances

In order to have moderate sized examples available for repetitive testing, we used subsets of the Dogrib landscape data. We generated two sub-instances that we labeled Sub-1 and

²http://www.firegrowthmodel.ca/prometheus/downloads/Dogrib_v624.zip

Sub-2 with a cell resolution of 100×100 meters. The first one represents a sub-forest from the Dogrib landscape that is 20×20 cells (400 ha) and the second a 40×40 cell (1,600 ha) instance. Both consist of heterogeneous landscapes that include different fuel types and non-flammable cells (e.g., mountains or rivers).

An ignition point was selected for each instance as a starting point for the fire growth validation experiments. Three weather stream files: Weather-1, Weather-2, and Weather-3 of 6, 14, and 22 hours respectively, were used as inputs. The first file contains data for the 6 hours during which the real Dogrib fire made a run (extreme weather conditions). The second and third files contain additional meteorological measurements from the same day of the fire, before and after that spread event. After the ignition point was fixed for both instances, we proceeded to run the simulation in Prometheus and the deterministic version of Cell2Fire, generated the hourly fire perimeters, and calculated the similarity metrics (1-MSE and SSIM).

British Columbia

The British Columbia instances set contains fuel data for five different areas – ArrowHead (265,536 ha), Revelstoke (391,314 ha), Mica Creek (348,404 ha), Glacier Natural Park (559,746 ha), and Central Kootenay (494,665 ha) – of the province. In order to validate our fire growth model, we defined two fires with random ignition points and 24-hour weather stream based on the historical weather dataset from the Climate Information Section of the Agriculture and Forestry site for each area and random weather streams generated for comparison purposes. These instances are provided with BurnP3³. We generated subsets of the large forests and simulated fire growth using Prometheus and Cell2Fire.

1.7 Computational Implementation

Analysis of the running times reveals that the initial ignition stage is negligible. The sending-messages stage updates the fire progress in every burning cell. Because a large number of cells can be burning at the same time and there are no direct dependencies on neighboring cells, updating the fire progress for each burning cell is easily parallelizable because the calculations for each cell can be done independently. Each cell, in addition to updating its current status, can also “send a burning message” to an adjacent cell. In the receiving messages stage, we analyze the “burn messages” sent to non-burning cells and mark them as burning if the fire start conditions are met. This part is also potentially parallelizable, but because the number of newly burned cells at a single time-step is dwarfed by the number of currently burning cells, we found that a speedup here is of lower priority.

Due to the easily parallelizable structure of our algorithm, the most suitable approach for parallelizing its execution consists of a shared-memory approach using the well-known

³http://www.firegrowthmodel.ca/burnp3/software_e.php

OpenMP API [103]. Using OpenMP is advantageous because the code is also optimized for execution on personal computers (see Appendix A for details).

1.8 Results and Discussion

We begin by describing experiments to validate the model. This is followed by illustration of the management capabilities and research issues.

Propagation Validation

In this section we compare the predicted burn perimeters produced by Cell2Fire and Prometheus for several hypothetical instances (described in Section 1.6) created for this purpose. We did not compare either simulator with the realized fire scars in our study because it is difficult to determine the extent to which the final shapes were influenced by suppression actions.

Sub-Instances

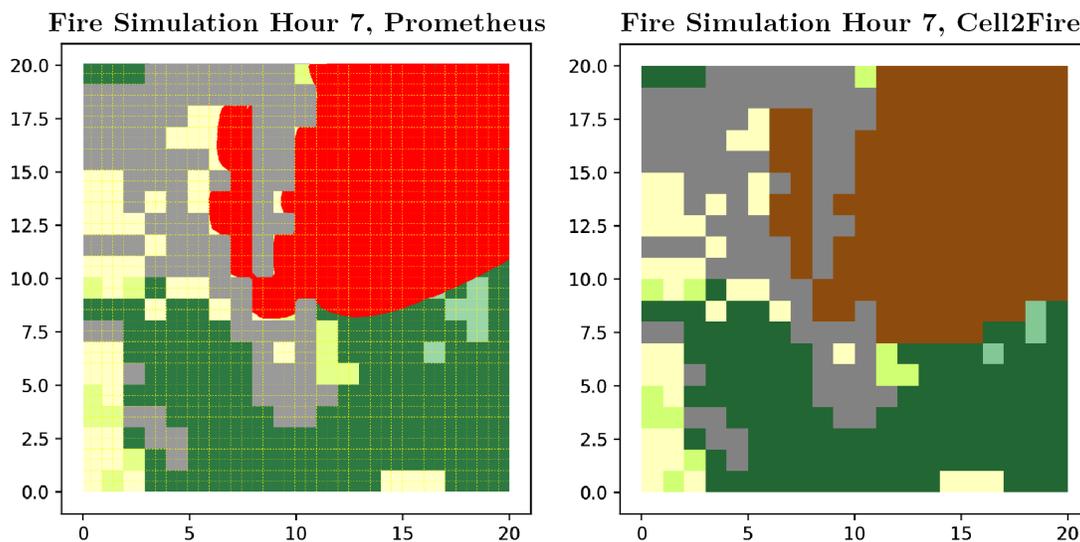


Figure 1.9: Fire growth visualization for sub-instance 2 using Prometheus and Cell2Fire across a heterogeneous landscape with non-flammable cells (mountains, gray cells) and different fuel-types (green and yellow cells).

Based on the results shown in Table 1.1 and Figure 1.9 we can see that Cell2Fire produced results that are similar to scars produced by Prometheus with respect to the hourly fire growth and final fire perimeter. When testing the 6 critical hours of the Dogrib fire, the level

of difference (MSE) is less than 5% for the first sub-instance and less than 6% for the second sub-instance. For the full day simulation, the differences are slightly larger, reaching average MSE levels values close to 6% and 7 % for the Sub-2 and the Dogrib instance, respectively.

On the other hand, we do sometimes observe a decrease in the structural similarity measurement (Table 1.1). The differences are more pronounced when using the 22 hour fire weather stream. These deviations indicate that the fire growth predicted by Cell2Fire differs from the one predicted by Prometheus due to: (1) the approximation of the elliptical model using a cellular-automata approach implies a different fire dynamic, increasing the differences in every time step compared to the wave front model and (2) differences in the calculations/approximations of the effective ROS values.

Hour	Sub-1		Sub-2	
	1 - MSE [%]	SSIM [%]	1 - MSE [%]	SSIM [%]
1	99.75	93.59	99.98	95.44
2	96.75	77.40	99.94	95.01
3	97.75	74.03	95.56	85.53
4	97.50	78.70	95.38	84.41
5	96.75	79.61	96.50	82.25
6	96.00	73.96	94.03	75.01
AVG [%]	97.42	79.55	96.90	86.25

Table 1.1: Accuracy measured by the complement of the mean squared error (MSE) and structural similarity index measure (SSIM) values per hour (6 hours evolution) for the sub-instances. Lower values in SSIM highlight the differences between the cellular-automata and wave propagation approaches as well as the definition of a burned cell between Cell2Fire and Prometheus.

The wave-propagation model based on the Huygens' principle implemented in Prometheus performs a series of approximations with respect to the burning area. Therefore, a cell is classified as a *Burned* cell in the Burn Grid output (0-1 Matrix) only if more than 50% of its area has been covered by fire (i.e., belongs to the interior of the approximated ellipse) whereas in Cell2Fire a cell is always either completely available or burned. In addition, an ignition point represents a complete cell in Cell2Fire (an area) while it is just a vertex/point for Prometheus, defining two different (but consistent) starting points for the fire spread evolution. This approximation based on discrete cells improves as the cell size decreases (i.e., data with higher resolution is available).

The hypothetical British Columbia wildfires

The final fire scars and similarity metrics, focusing on the affected area of the instance for easier visualization, obtained for the 10 simulated wildfires using both Prometheus (columns

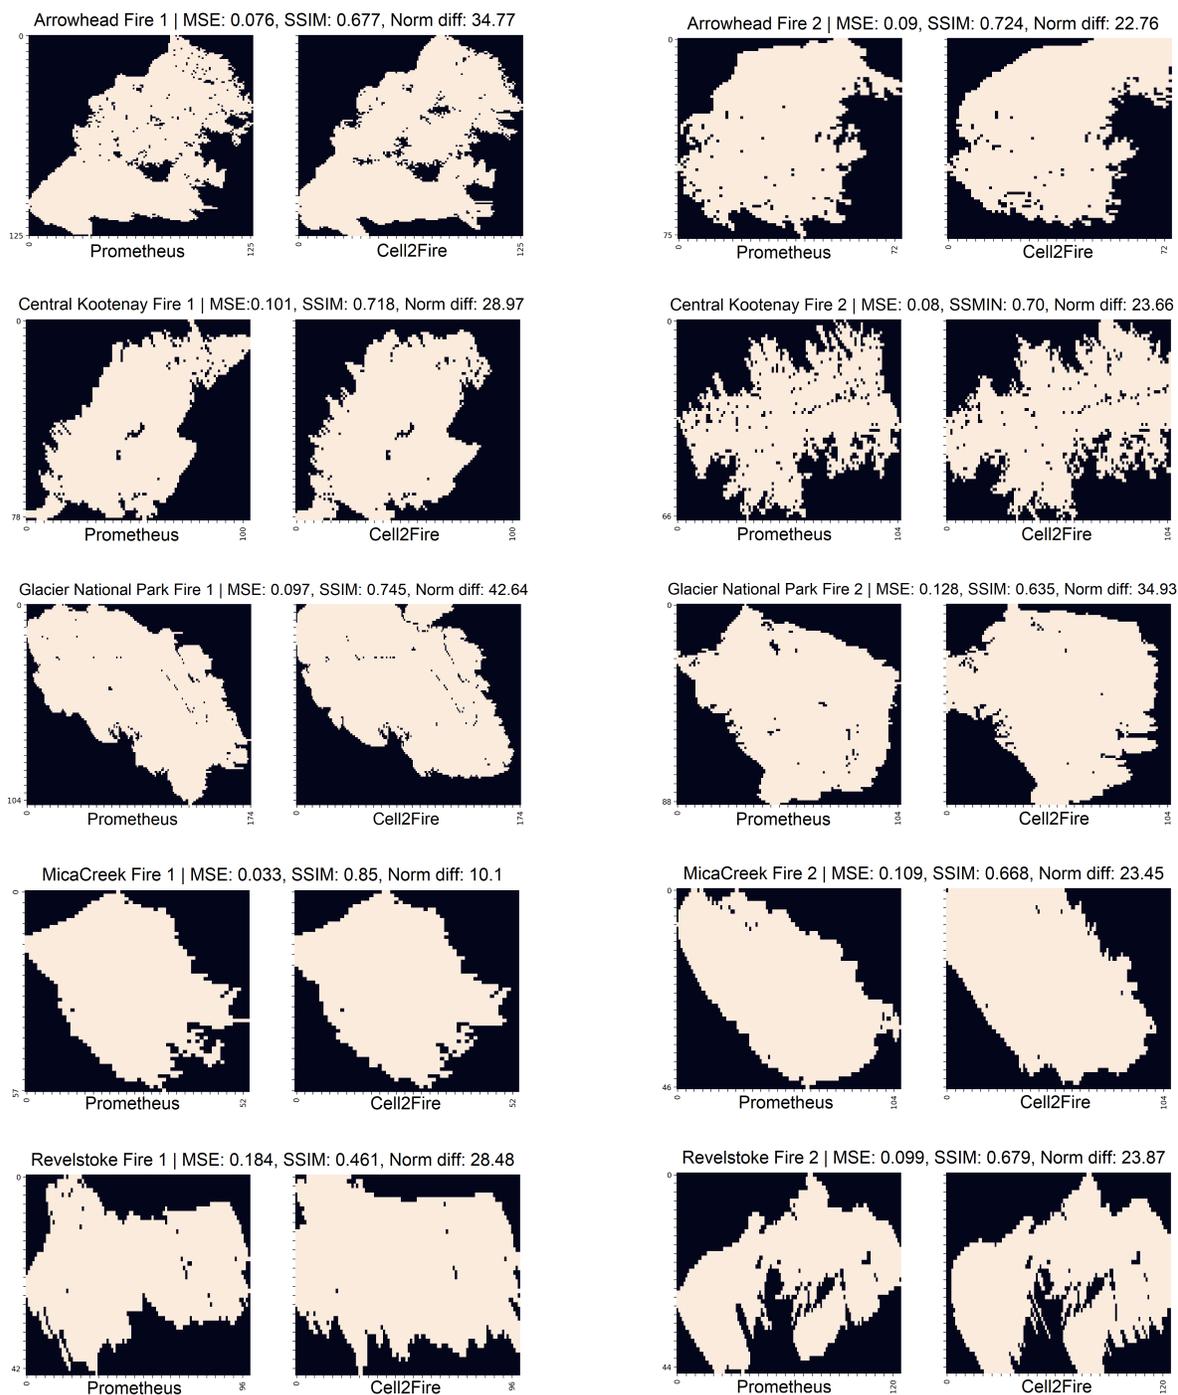


Figure 1.10: British Columbia wildfire instances. From the final scars and statistics (left side Prometheus, right side Cell2Fire), we see how the Cell2Fire scars compared to scars produced by Prometheus, reaching $\overline{MSE} = 0.0995$, $\overline{SSIM} = 0.6863$, $\bar{\delta}_{norm} = 16.347$.

1 and 3) and Cell2Fire (columns 2 and 4) can be seen in Figure 1.10. Those results exhibit a high similarity between the scars across the main three metrics for all forests, as demonstrated in Table 1.2. From the results, we note an average MSE of 9%, with a maximum MSE and minimum SSIM of only 18% and 46%, respectively (Revelstoke 1 instance). High-performance average SSIM and δ_{norm} values of 68% and 27.4% are obtained across all forests, respectively, reflected in very similar final perimeters in all landscapes. The best results are obtained in the MicaCreek 1 instance, with a 3.3% MSE and an 85% of SSIM, translated in almost identical wildfires. We note that, in general, Cell2Fire tends to burn more cells than Prometheus. As previously noted, this pattern was expected due to the different approximations used by both approaches to defining a burned cell as well as the ignition point/area.

	<i>MSE</i>	<i>SSIM</i>	δ_{norm}
Mean	0.09	0.68	27.36
Std	0.04	0.09	8.88
Max	0.18	0.85	42.64
Min	0.03	0.46	10.10

Table 1.2: British Columbia simulations summary statistics obtained by comparing the simulated final fire perimeters from Prometheus and Cell2Fire. Mean squared error (MSE) and structural similarity index (SSIM) are reported across all instances. We observe low $\delta_{norm} = ||X - Y||$ where X and Y are the binary BurnGrids matrices produced by both simulators, indicating the high similarity between the results of both simulation approaches.

These results indicate that Cell2Fire approximates the results of a state-of-the-art simulator like Prometheus for different fuel types, landscapes, and weather scenarios. Different ignition points and weather scenarios were tested on these landscapes, obtaining similar results in terms of the main similarity metrics. A similar pattern was observed with respect to the hourly evolution of the fire perimeters.

Dogrib fire instance

The comparison of the similarity metrics after 22 hours of fire growth is shown in Figure S2 and Table S2, where both $(1 - MSE)$ and *SSIM* values can be seen for each simulated hour. We observe that Cell2Fire predicted a very similar growth with respect to the Prometheus wave-front approach, obtaining good performance when compared with Prometheus, not exceeding a 12% and 18% of difference in both measurements. An average of 87.91% of structural similarity and a global average of 91.82% of accuracy ($1 - MSE$) were obtained during the 22 hours of active fire growth. A clear pattern can be seen where both similarity metrics start high and remain stable during the first 4 hours, then they show a significant

negative slope between hours 4 to 11, and finally reach a steady state for the rest of the simulation. The explanation behind these results is clear: during the initial 4 hours of the fire, similar fire growth occurs due to weather conditions that are not extreme; however, weather factors between hours 4 to 11 exhibit the most extreme conditions (strong wind speed, high temperatures, etc.) magnifying the fire growth differences/approximations between both approaches in terms of the number of burning cells per hour (fire scar). After hour 11, differences between fire scars tend to be stable due to the lack of new extreme weather episodes. However, the structural differences remain in the fire perimeters obtained in the previous hours.

In addition, some of the structural differences between the generated fire scars can be explained in part by modeling features included in Prometheus but not in Cell2Fire such as *Breaching*, where non-fuel grid cells or linear fuel breaks fail to stop an advancing fire front, a feature that is not currently included in Cell2Fire. We conclude that Cell2Fire produces results that are similar to the ones produced by Prometheus. The final fire scars are also similar as seen in Figure 1.11, where the simulated fires (left and right) and the real satellite (center) images are shown.

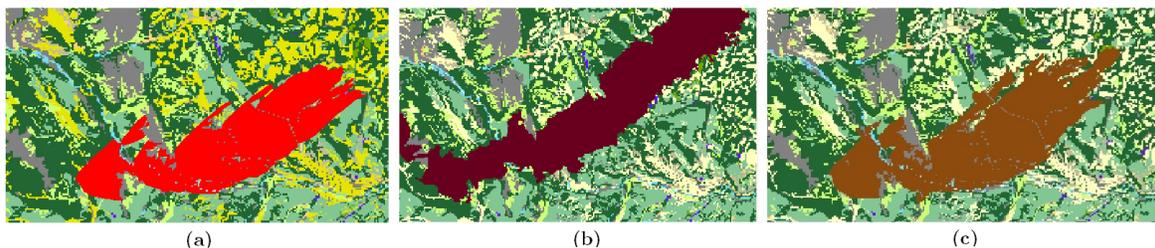


Figure 1.11: a) Prometheus fire scar obtained for the region of Dogrib fire area, contrasted with the b) real fire projected into grid format in 2002 and c) Cell2Fire final output.

Managing the landscape

To illustrate use of the management module we focus on the evaluation of harvesting plans in the 1,600 ha landscape (sub-instance 2) located in the province of Alberta (see Figure 1.12). This area is characterized by a fragmented land cover distribution, including the same fuel types as in the Dogrib instance (grass, conifers, and non-flammable fuels). For evaluation purposes, we set the ignition cell at random at each replication and simulate multiple wildfires using representative weather scenarios obtained from the closest weather stations for a total of 12 hours. A second set of experiments using a fixed ignition point is also included for comparison purposes. We compare the outputs and wildfire behavior variables obtained before and after applying three different harvesting schedules to 10% of the available cells: (1) baseline approach harvesting cells completely at random; (2) focus

the harvesting on cells with higher burn probability estimated from 100 simulations on the original landscape; and (3) a myopic plan focusing its attention on those cells with higher economic value. We explicitly incorporate wildfire uncertainty running 100 independent replications for the pre and post-treatment landscapes.

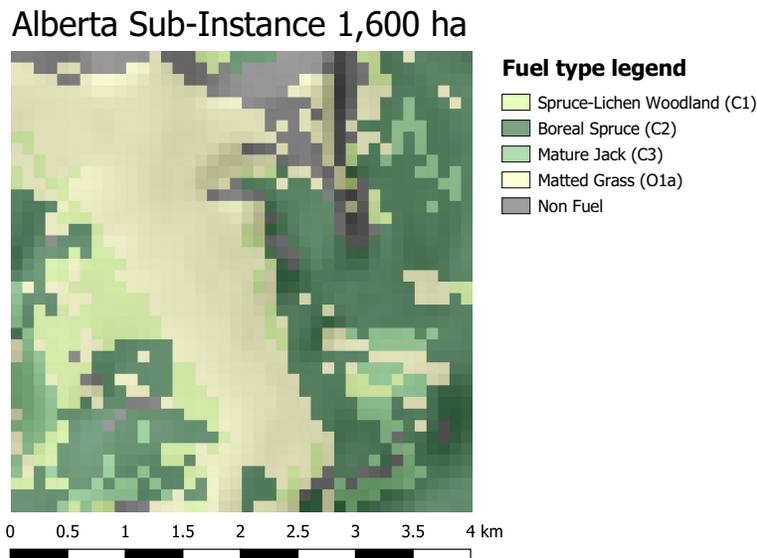


Figure 1.12: Map of the Alberta Sub-instance. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of the area. A hillshade effect has been applied to depict the elevation of the terrain.

Multiple outputs are generated after the simulations in order to evaluate the performance of the management plans (Figures 1.13 and 1.14). These outputs include burn probability maps representing the wildfire susceptibility of each cell under the ignition and weather conditions used; ROS heatmaps indicating the average rate of spread over the simulated period as well as highlighting areas of the landscape experiencing high/low propagation rates; a consolidated shortest-path tree (directed network flow graph representing the fire spread dynamics where nodes are the cells of the landscape and edges indicate fire propagation statistics between them) representing the critical, fastest, and most frequent propagation paths of the wildfire in the landscape; and multiple plots and statistics comparing the performance of pre-treatment and post-treatment landscapes given a performance metric such as expected area burned.

From the outputs of the simulations with random weather (RW) on the original pre-management landscape using random (RI) and deterministic ignition (DI) points (Figure 1.13), we observe areas with higher burn probability values on the top right section of the landscape according to the historical wind direction distribution (northeast). This pattern is exacerbated in the fixed ignition experiments. In this case, we observe that wildfires are

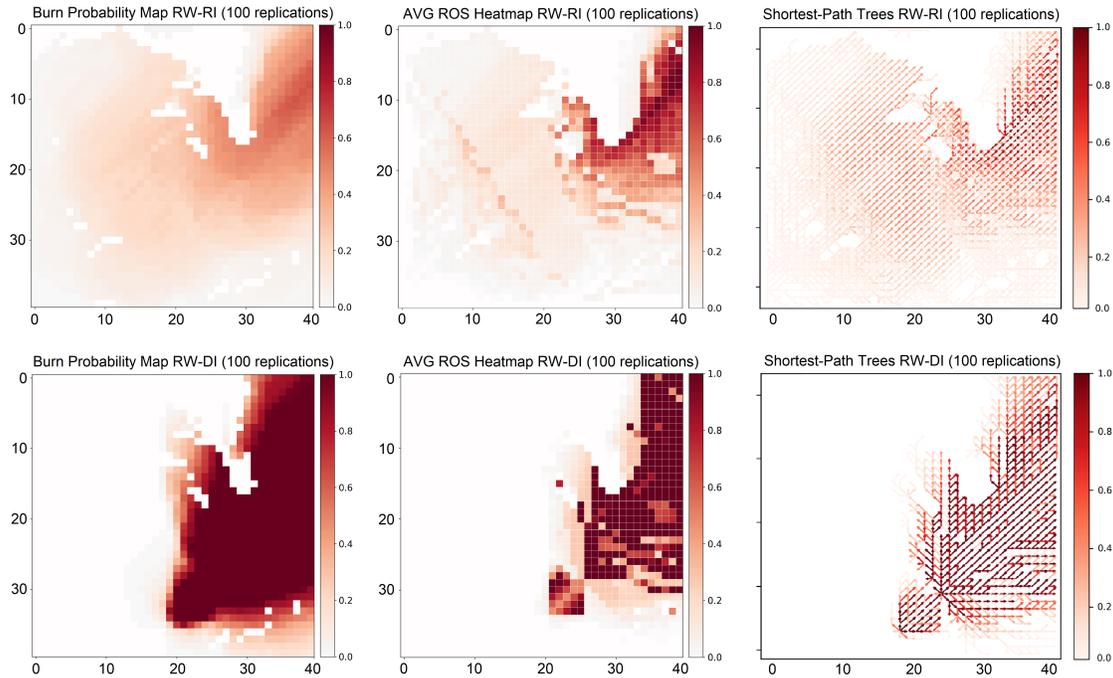


Figure 1.13: Pre-treatment analysis. Burn probability maps (left), average ROS heatmaps (center), and consolidated shortest-tree paths (right) are generated for the random weather-random ignition (RW-RI) and random weather-deterministic ignition (RW-DI) experiments. Darker areas indicate higher values (normalized). Fixing the ignition point (RW-DI) is translated into a very focused wildfire dynamic on the eastern side of the landscape, with no fires escaping to the western region of the instance. This results in higher ROS average values in several cells on the east given the dominant land cover in the area and a denser shortest-path tree than the one observed in the random ignition (RW-RI) experiments since all fires are condensed in a unique side of the land.

not able to reach the western side of the landscape, focusing the spread on the eastern and north-eastern areas. In addition, we note the non-flammable section of the landscape at the top center area of the instance characterized by a chain of mountains, preventing the spread of the fire. Observing the ROS heatmaps, we notice that wildfires reach their maximum ROS on the areas covered by conifers (Boreal Spruce) obtaining average ROS values close to 5 [m/min] and 30 [m/min]; and peak values of 34 [m/min] and 118 [m/min] for the RW-RI and RW-DI experiments, respectively. Finally, the consolidated shortest-path trees highlight the areas where fire tends to propagate (darker zones) observing how the area covered by conifers tends to be actively involved in spreading fire to adjacent regions.

Focusing on the most interesting RW-RI experiments (Figure 1.14), we observe how these outputs are perturbed after applying different harvesting plans. In the case of the

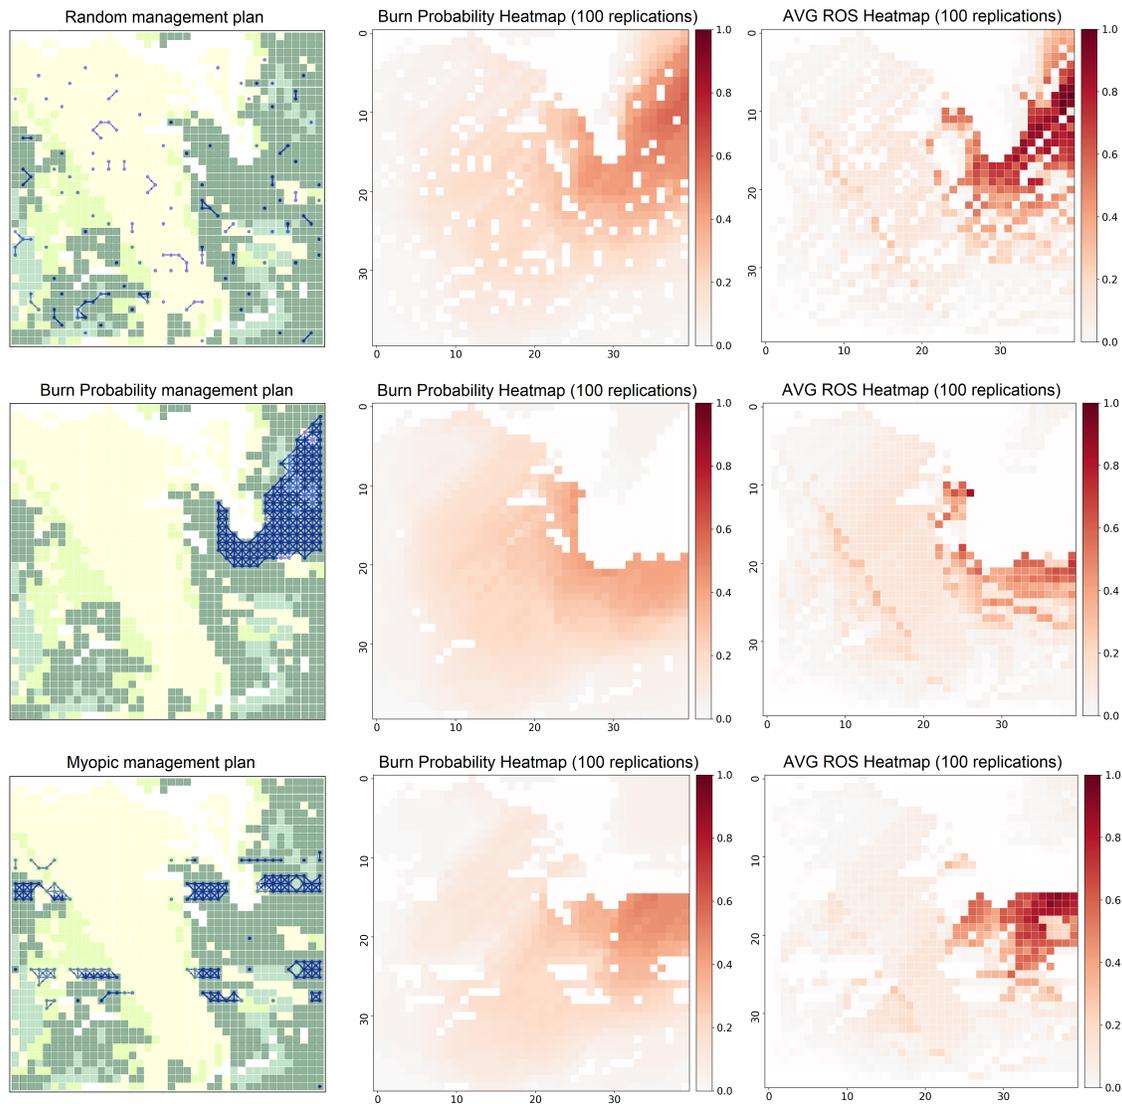


Figure 1.14: Post-treatment RW-RI experiments. Management plans (left) are generated using different policies to select which cells to harvest (blue dots). Non-flammable cells are represented by white cells for visualization purposes. Burn probability maps (center) and average ROS heatmaps (right) are generated from 100 independent replications to evaluate and compare different metrics for the selected treatment plans.

random management approach (top row), the uncorrelated fragmentation introduced in the landscape is not an effective measure to reduce fire spread. This is reflected in the expected area burned, which is reduced by only 22.3% compared to the non-treated landscape. In addition, the average ROS is reduced by 17.5% but reaches significant peaks (32.9 [m/min])

in the northeast areas. Moreover, it is not practical (e.g., to move the harvesting equipment across the whole landscape “at random”) and clearly not profitable for the landscape manager. The BP-based plan (center row) improves the impact on the expected area burned, decreasing it by 40%. In addition, it effectively reduces the average ROS by a 37% compared to the non-managed land (with peaks up to 24.5 [m/min]). Still, it may not produce a feasible plan from an economic perspective as it does not consider the revenue layer when selecting the cells to harvest. Finally, we observe how the myopic approach (bottom row) that focuses the harvesting plan on the areas where the most profitable cells are located. However, this may become a limitation of the plan as it does not incorporate wildfire behavior features when selecting the treated areas. Although it decreases the expected area burned by 31%, its limitations are reflected in the limited impact on the peak ROS (30.1 [m/min]) only decreasing it by 10% compared to the pre-treatment state. On the other hand, it has a significant impact on the average ROS 3.1 [m/min], decreasing it by a 38%. This is explained because, in this instance, the most profitable cells match the land covers experiencing higher ROS during the simulations, thus allowing the treatment to be effective in reducing the average ROS values. However, this management plan results in larger variance than the BP-based treatment, exposing the landscape to higher risk levels. Using this framework, planners can easily compare some of the most relevant outputs and evaluate alternative treatment plans in an effective and efficient way (less than one minute) to develop robust and objective management plans.

In addition, given the graph structure of the wildfire propagation patterns generated by the simulator, we can apply a series of complex network algorithms on top of the generated outputs (shortest-path trees) to evaluate the wildfire behavior in the original landscape and the impact of the management plan. As an example, decision makers could calculate metrics such as the average betweenness centrality (BC) [63] of each cell in the landscape across all replications to identify which nodes have a more active role in the propagation of wildfire to other parts of the land (Figure 1.15). Alternatively, a degree heatmap indicating the average outgoing degree of each node can be generated, among several others useful metrics. Using this information, managers could decide to modify their initial plans to focus on those critical areas where wildfire tends to propagate faster and more frequently.

Observing the BC-based treatments for the RW-RI and RW-DI experiments (Figure 1.15), we notice how nodes located in the northeast section of the landscape are selected. They tend to be frequently involved in the fire propagation dynamics, potentially playing a fundamental role to mitigate future expected losses due to large wildfire events in the area. Comparing the pre- and post-treatment simulation results, we observe how the distribution of both heatmaps is impacted: shifting the areas with highest/lowest weights across the landscape; reducing average ROS values by 42% (RW-RI) and 46% (RW-DI); and decreasing the expected area burned by 36% (RW-RI) and 93% (RW-DI). This information could play a crucial role in selecting which areas to allocate future wildfire suppression resources, focusing the efforts on those unprotected or fire-prone areas detected from the generated outputs.

Therefore, we observe how Cell2Fire can be a valuable tool to assist decision-makers when defining management plans. It helps them understand and generates relevant managerial

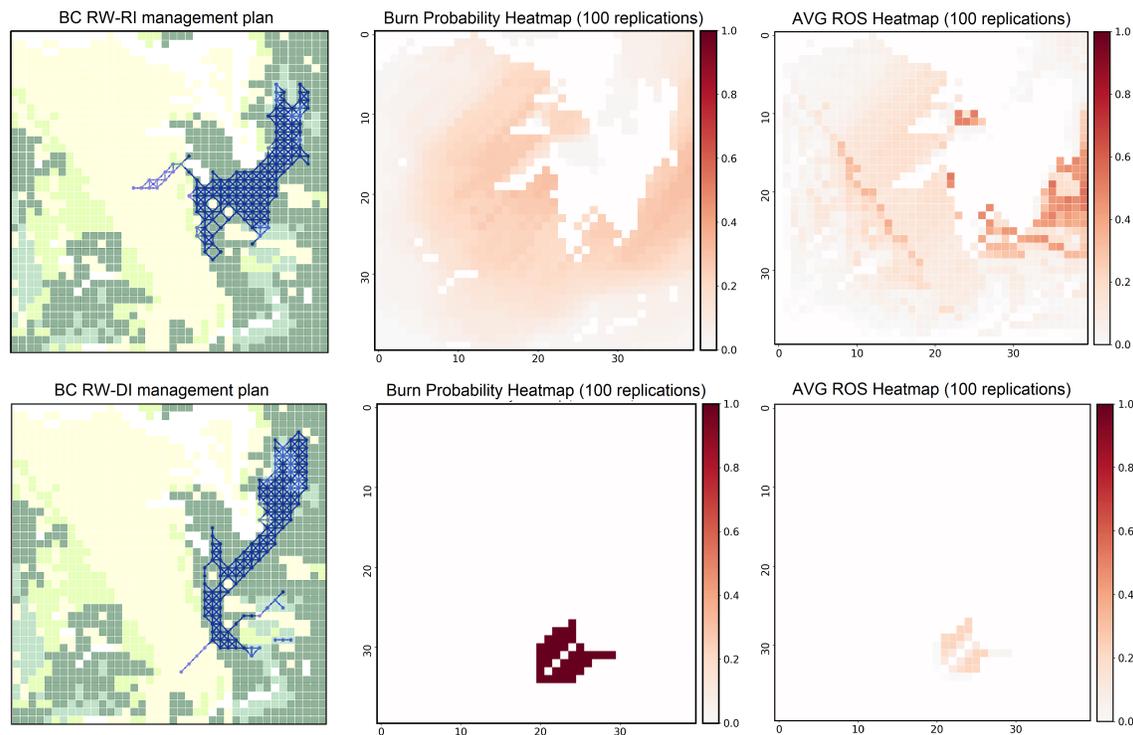


Figure 1.15: Post-treatment BC plans. Management plans (left) for the RW-RI and RW-DI experiments are generated using BC as the main metric to select which cells to harvest (represented with blue dots). Non-flammable cells are represented by white cells for visualization purposes. Burn probability maps (center) and average ROS heatmaps (right) are generated from 100 independent replications to evaluate and compare different metrics for the selected treatment plans.

insights regarding which areas of the landscape are more/less prone to future wildfires, the most likely areas to which the fire can propagate after ignition in the treated-landscape, the expected rate of spread/burned area across the available cells, among other relevant information.

Perhaps more important, is the support that Cell2Fire can provide to researchers studying methods for forest management planning (practical examples for the interested reader can be found in [240, 252]). All this, to support and improve the whole decision-making process, providing the opportunity to analyze the complex and significant trade-offs involved in landscape management under wildfire uncertainty.

1.9 Conclusions

Cell2Fire provides numerous opportunities for researchers interested in incorporating fire growth in their models for strategic harvest planning and fuel management planning. Researchers can easily test and adjust their decision-making models to enhance their management plans and identify relevant metrics to capture and actions to mitigate the impact of wildfire uncertainty. In addition, it incorporates a series of out-of-the-box planning heuristics that provide baselines to the decision-makers and a series of complex network algorithms to evaluate the wildfire behavior in the original landscape and the impact of the management plan. This information could play a crucial role in selecting which areas to allocate future wildfire suppression resources, focusing the efforts on those unprotected or fire-prone areas detected from the generated outputs. We will expand on the decision-making aspects in Chapters 3-5, where we will formalize the landscape planning problem as an explicit optimization problem.

Using the FBP fire spread model, we have compared the simulated fire perimeters with perimeters produced by a state-of-the-art simulator for validation purposes. Other fire spread models can be employed instead, which extends the range of environments where Cell2Fire can be used as we will cover in Chapters 2 and 4. In addition to supporting stochastic ignition and weather, the simulator also supports random sampling of the ROS. Stochastics are a major area of further research. We also plan to initiate research concerning slow fire growth at night as well as spotting which allows fires to jump across unburnable barriers.

Because the software is open-source and modular, it lends itself to customization as needed. The simulator is fast and it scales well in parallel computing environments, so it is well-suited for use with large, flammable landscapes and in studies that require many simulations. By adding a highly parallelizable, open-source fire growth simulator to the tool set available, we hope to provide transparent support for ongoing research.

Chapter 2

Learning from historical data: a general fire spread model using derivative-free optimization

2.1 Introduction

In Chapter 1, we have studied the problem of simulating wildfires accurately and efficiently while integrating decision-making models for landscape planning under wildfire uncertainty. We note that an independent fire behavior spread model (e.g., the Canadian FBP or BEHAVE in the U.S.) is required to accurately simulate wildfires in a determined region. These models, traditionally developed via long-term experimental efforts performed in controlled environments (inside laboratories and/or on the field) are needed to approximate the ROS of the wildfire. This computation allows us to simulate fire growth for specific time windows and conditions.

However, such models are not generally available to use in alternative regions. This, due to modeling assumptions leading to inaccurate results when implemented in different conditions to the ones for which the model was developed, or due to limited resources and data access in the region of interest. Therefore, one of the main limitations of these models is that they cannot be easily applied to regions with different characteristics, as several assumptions need to be satisfied: land covers must match the ones incorporated in the original model, all inputs must be available in the same format as in the region of origin, among several other restrictions, thus, limiting fire behavior models to specific territories. Moreover, these models are mainly static, in the sense that they were developed under specific environmental conditions that could be significantly different from the current ones (e.g., due to climate change), potentially leading to inaccurate results if no adjustments are performed.

To address this limitation, in this Chapter, we explore the creation of a general fire spread model with the capacity of exploiting historical data to automatically adjust the

main parameters of the fire spread model using a derivative-free optimization approach. This, to obtain more accurate, flexible, and transferable fire spread models allowing us to effectively simulate wildfires in any region of the world where historical data is available.

Considerable efforts have been made in recent decades to understand the dynamics of forest fires on landscapes in different parts of the world due to their increase and their significant effects on our ecological and human systems [290]. A series of relevant research studies focused on understanding fire behavior have determined that fire is mainly influenced and directed at the landscape level by four main factors: i) weather, ii) topography, iii) humidity of fine vegetable fuels, and iv) type of forest fuel [10, 156, 267]. In this context, it is known that high temperatures and strong wind gusts can produce fires of great magnitude and that are difficult to control [359], with the potential of affecting multiple components of the ecosystem. On the other hand, terrain slopes affect the rate of spread of the fire, increasing and decreasing it with uphill and downhill slopes, respectively. The elevation of the terrain could influence the humidity of the leaf of the trees and shrubs, and certain bands of elevation can sustain climatic and environmental conditions that are more/less fire-prone such as dry and windy flammable areas, or wet and cold rocky regions. Finally, the type of forest fuel (land cover) that goes into combustion may have different flammability and therefore, release different levels of energy. This could lead to multiple damage levels to ecosystems and become a significant menace to our human communities.

It is because of the aforementioned, that scientists in different countries have developed or adapted complex systems to understand this phenomenon. Fire behavior models have been under development since [287] and [353] works, with actual applications in wildfire control, risk assessment, prescription development, long-term planning, and in first and second-order fire effect models at different spatial and temporal scales [277]. However, the basics for any fire behavior model application are the accurate prediction of the burned area and fire intensity of a simulated wildfire. A precise prediction of these two characteristics of a wildfire needs complete knowledge of the relationship between fuel and the atmospheric conditions as well as its interaction with the topography in local landscapes. For decades, a few countries have invested in the science needed for this kind of knowledge, mainly restricted to the U.S., Australia, Canada, Spain, and Portugal [69, 156, 322]. For example, in Canada, the Fire Behavior Prediction (FBP) System [156] is the system dedicated to modeling fire behavior for static conditions, mainly based on controlled experiments in laboratories to measure and relate the different variables mentioned that influence the multiple aspects of fire [13, 156], as discussed in the previous Chapter. Its analog in the United States is the Fire Behavior Prediction and Fuel Modeling (BEHAVE) System [69]. Another local fire behavior system and fire growth simulator that models the behavior of fire at the surface level is Kitral, developed in Chile in 1995 [170]. Although being a simpler system than the ones implemented in the northern hemisphere and Australia, it has shown to fit quite well with the reality of local fires [265]. A complete review of the main fire behavior modeling systems can be found in [130].

Despite these efforts, even in those countries with state-of-the-art systems, inaccuracies presented in fire behavior models are still an open research problem because of the high

degree of uncertainty in the classification of fuel types, the estimation of the actual fuel moisture, the quality, and accessibility to local atmospheric data, or even the specific fire behavior model parameters used depending on the observed conditions [177]. Challenges such as i) scaling-up existing fire behavior models from point-scale to hundreds of kilometers of wildfires without underestimating the local effects of landscape spatial patterns, ii) integrate and take advantage of the data revolution into fire behavior models, iii) obtain a better and comprehensive understanding of the fuel-fire-atmospheric relationship in real landscapes, and iv) explicitly incorporate the human interaction with fires in management and control [222] are some of the most relevant research opportunities in the field that should be addressed by the scientific community to be able to tackle the current complex global scenario.

In the same line, a wide set of fire growth simulators are available from multiple regions, which utilize a range of different modeling approaches and underlying fire behavior prediction systems to simulate the fire spread dynamics based on demographics, topographic, and environmental conditions [256]. To our knowledge, the most used simulators are based on rate of spread models derived from the fire behavior systems mentioned above, which were predominantly developed using observations of experimental fires [322]. However, the conditions under which wildfires occur are not easily replicated experimentally because of the costs and dangers that this entails. The latter could lead to significant differences in the behavior of the observed fires and the simulated ones. Moreover, existing systems may not take advantage of historical data in automated (e.g., online learning) ways due to the complexity of incorporating data-driven learning models in the simulation framework to adjust the original parameters of the system or due to the high investment of resources (time, experiments, etc.) and the need of experts' knowledge to adapt the existing models according to the current data. Therefore, this challenging setting could be translated into a significant development barrier to adapt existing fire behavior models to other regions as well as the presence of more static systems, with difficulties to account for new and evolving uncertain climatic scenarios and complex propagation patterns observed in real events. Thus, the applicability of the models and potentially, their accuracy, could be impacted if no significant resources are invested to maintain and adjust state-of-the-art fire spread models in the long-term.

The systems mentioned were the results of innumerable research studies across many years, supporting their performance with verifiable evidence through historical events analysis. This fact, highly positive on the one hand, represents a serious obstacle for less developed countries that lack or present scarce basic information for the correct transfer of the systems to their present reality and, even less, the ability to properly adjust them to their own environmental and geographical conditions. The adaptation of existing fire behavior models or the development of new and flexible ones to face these challenges arises as a global issue for emerging countries: the technical and scientific knowledge needed for the regular use of fire behavior models for multiple applications and the number of required resources, act as a barrier to their adaptation and potential implementation in different regions, being even more acute in less developed countries. All this, exacerbated with increasingly extreme weather conditions resulting from climate change where wildfire occurrence is likely to con-

tinue to extend well beyond those conditions under which existing fire behavior models were developed [290, 359, 360].

As we have discussed in the previous Chapter, capturing and accurately predict fire propagation patterns could play a crucial role to support decision-makers. In the context of landscape management under wildfire uncertainty, it is urgent to derive effective prevention and mitigation policies to address the impact of wildfires on human well-being, the conservation of biodiversity, and their effects on greenhouse gas emission, among many others [61, 176]. Therefore, the development of a flexible and adjustable fire spread system, with the ability to learn and extract propagation patterns from historical data remains a relevant and open research problem, being the main motivation of our research.

In this Chapter, we focus on the development of such a model, adjusting it for different regions and testing its integration with a fire growth simulator. The objectives of this study are four: i) propose a general fire propagation model that considers the main aspects of the most specialized fire behavior systems currently developed; ii) introduce a methodology based on derivative-free optimization algorithms integrated into a fire growth simulator to adjust relevant parameters of the fuel models automatically through historical fires in different areas of the world; iii) decrease the gap between theory and practical applications in the context of wildfire simulation; and iv) discuss the scope, limitations, and future challenges of this methodology to develop fire behavior systems in countries where no spread models are available for wildfire simulation.

This Chapter is organized as follows. In Section 2 we introduce the theoretical background and fundamental components of our proposed general fire spread model that will be adjusted with historical data to mimic observed fire events. Section 3 presents our proposed data pipeline and processing algorithms allowing an automated extraction of historical wildfire perimeters for training the spread model. Section 4 introduces the mathematical models and describe the derivative-free algorithms applied to solve them to adjust the main parameters of the general fire spread model. Section 5 presents the computational results of our experiments in three regions illustrating the potential of the proposed methodology. Section 6, contains our conclusions and thoughts concerning future research needs. Appendix B provides details concerning extensions of the fire spread model and its computational implementation.

2.2 A General Fire Spread Model

In order to develop a general fire spread model, we require to capture the main characteristics of a forest fire: i) define a surface propagation model that reflects the way fire spreads through a landscape, ii) calculate the rate of spread (ROS) in multiple spread directions to predict its potential evolution, and iii) the inclusion of variability under multiple climatic conditions to obtain robust estimates. As a starting point, we use the relations included in the most complete fire behavior systems to date, which to our knowledge are the American [69], Canadian [156], and Australian [322] systems. Current systems model the main dynamics of fire behavior under certain homogeneity assumptions: fuels and topography are uniform and

continuous and the wind direction is constant and unidirectional. This step is replicated for all land-cover categories included in the fire spread model, obtaining independent parameters for each fuel type that best mimic the experimental data. Then, these dynamics (i.e., the parameters of the models) are adjusted by factors accounting for heterogeneity such as slope variations, wind fields, and moisture levels. In this section, we only refer to the key aspects to model the dynamics of fire, as a function of time through a heterogeneous land. Details of secondary components of the spread model are presented in Appendix B.

Surface propagation model

Similar to the fire behavior model discussed in the previous Chapter, all the mentioned systems consider an elliptical propagation model over a two-dimensional lattice to model wildfire spread. In these models, each cell represents a homogeneous land cover unit and is characterized by specific topographic/demographic characteristics features such as elevation or slope (Figure 2.1). The dimensions and shape of the ellipse are estimated as a function of the land-cover and environmental conditions, being the wind speed (WS) one of the main disturbance factors. This perturbation is calculated and reflected in the fire propagation dynamics using the length-to-breadth ratio (LB) of the ellipse, defined as the simple quotient between the major and minor axes of the ellipse. Therefore, ellipses with large LB values tend to be associated with higher WS values, i.e., faster propagation patterns in the main axis of the ellipse (long-shaped ellipses) while lower LB values represent slower propagation patterns where the fire tends to spread in all directions with the same ROS, e.g., , in the extreme case of $LB = 1$ where fire propagates following a circular pattern (Figure 2.2).

In this setting, the fire expansion is calculated for different directions from the focus of the ellipse, located at or near the ignition point [14, 156]. A propagation template, defining which cells will be considered adjacent cells for keeping track of the fire progress (e.g., eight cells as in Fig. 2.1), is included as part of the fire growth model. These templates can differ in two main dimensions depending on the system: 1) they can be symmetric or asymmetric, meaning that cells in all directions with respect to the burning one are considered, or specific orientations are given more weight during the propagation dynamics; and (2) the number of cells considered. Then, the fire progress is updated between adjacent cells by using the outputs obtained from an independent fire spread model (e.g., FBP). However, different systems have found different equations that best fit the conditions of their own regions. For example, the FARSITE simulator in the U.S. uses the following equation to calculate LB :

$$LB(WS) = 0.936e^{0.9238 \times WS} + 0.461e^{-0.5573 \times WS} - 0.397 \quad (2.1)$$

while the BEHAVE system considers

$$LB(WS) = (0.936e^{0.9238 \times WS} + 0.461e^{-0.5573 \times WS} - 0.397)^{0.46} \quad (2.2)$$

On the other hand, in Canada, the Prometheus simulator and FBP System consider Eq. (2.3) for the majority of forest fuels based on empirical data, except for grass fuel types

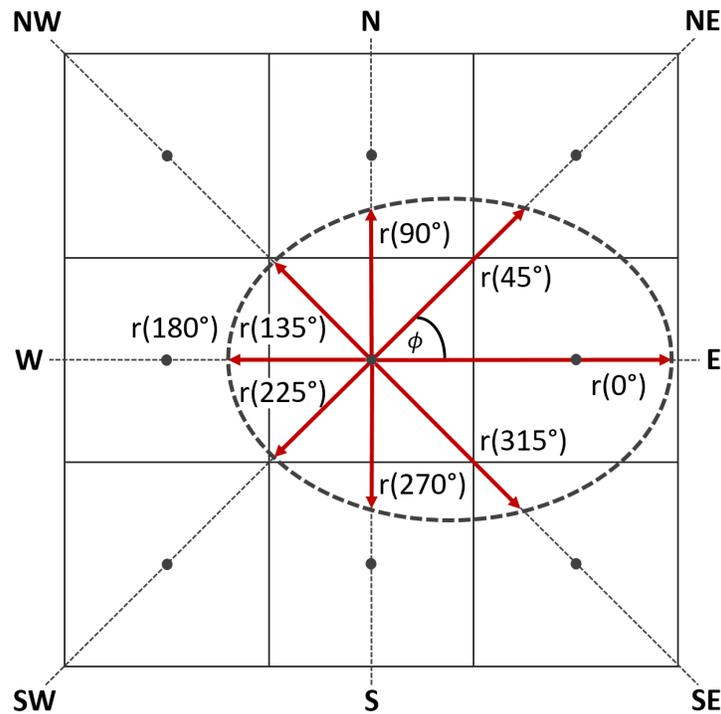


Figure 2.1: Surface propagation model elliptical expansion. An example of the elliptical propagation model is shown when the main wind direction aims to the east ($WD = 0^\circ$) and the focus of the ellipse is located at the center of the middle cell. $r(\phi)$ is the estimated rate of spread (ROS) in the ϕ direction with respect to the main wind direction WD .

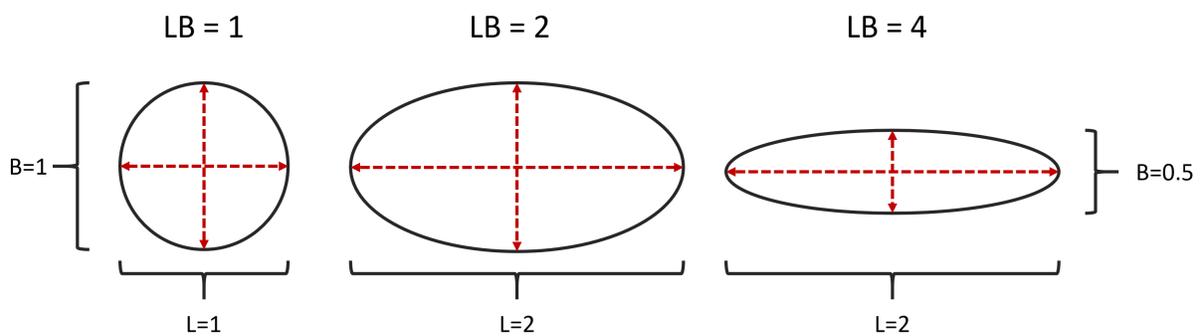


Figure 2.2: Length-to-breath ratio and ellipses. Examples of ellipses with different LB values are shown, visualizing different potential propagation patterns depending on the shape of the ellipse.

denoted by O-1 where Eq. (2.4) is applied.

$$LB(WS) = 1 + 8.729 [1 - \exp(-0.03 \times WS)]^{2.155} \quad (2.3)$$

$$LB(WS) = \begin{cases} 1.1 + WS^{0.464}, & WS \geq 1.0, \\ 1.0, & WS < 1.0. \end{cases} \quad (2.4)$$

It should be noted that this last equation (Eq. (2.4)) was taken from McArthur’s experimental studies of forest fires in Australia [219], which reveals that it is a common practice for different fire behavior systems to exchange useful knowledge and adapt it to their regions.

From all the equations, we observe that, as expected, higher values of WS are translated into long-shaped ellipses (i.e., larger LB) while less windy conditions are reflected in wider ellipses, representing slower fire propagation dynamics. Moreover, the expressions reflect how the proportions of the ellipses are only disturbed (usually as an exponential factor) by WS in these systems, being a fundamental input. These proportions are then applied for the different fuel types included in the fire spread model, modifying the dimensions of the ellipses according to their characteristics (e.g., flammability) while satisfying the LB equations.

Different is the case of the Kitral system, developed in Chile [170]. It does not explicitly rely on elliptical propagation. Rather, it is based on what the authors call “propagation factors due to wind speed” for directions other than the main direction (WD).

For this reason, based on the experimental data reported by Julio (1993) in [170], we propose the following general expression for the LB , as a function of wind speed (WS) in [Km/h] and two main adjustable shape parameters l_1 and l_2 :

$$LB(WS) = 1.0 + [l_1(1 - \exp(-l_2 \cdot WS))]^2 \quad (2.5)$$

These two shape parameters (l_1 and l_2) provide us with enough flexibility to represent all previous LB equations by simply adjusting their values, allowing us to reproduce the results of these validated fire spread models. For this, the optimal $\vec{l} = (l_1, l_2)$ values minimizing the approximation error and maximizing the coefficient of determination (R^2) compared to the original models, are calculated with the MATLAB r2018a Curve Fitting Toolbox using the nonlinear least squares (NLS) method [214]. These values are fitted using the outputs from the mentioned systems, where multiple LB values are calculated for different WS levels up to 60 [km/h] to obtain data points for the fitting procedure. Using the parameters found by the fitting method, we can determine, for each system, the corresponding l_1 and l_2 values that best represent the LB curves modeled in each system, as shown in Table 2.1 and Figure 2.3.

The importance of the previous equation is twofold, thinking that we need a general model for the LB . First, our results show that it has high flexibility to adapt to the four studied models, reflected in the high R^2 coefficients obtained with an average value of 0.99 (Table 2.1). Thus, we can obtain almost identical LB values to the ones used in these models. Second, these values will be relevant to be used as starting points for the automatic adjustment procedure. They provide us with low-error initial values, thus, improving the convergence of the proposed data-driven optimization framework.

Table 2.1: Elliptical length-to-breadth fitted shape parameters l_1 and l_2 for the proposed general LB expression. Optimal values are obtained for FARSITE, FBP, Kitral, and BehavePlus systems using a non-linear fitting procedure. The coefficient of determination (R^2) is used as the main performance metric for the fitting.

System \ Param.	l_1	l_2	R^2
FBP System (Others)	3.053	0.02667	0.999
FBP System (Grass)	2.454	0.07154	0.969
Anderson (1983) - dense forest stand	1.411	0.01745	0.993
Anderson (1983) - open forest stand	2.587	0.01142	0.995
Anderson (1983) - grass/slash	5.578	0.006023	0.996
Anderson (1983) - heavy slash	37.49	0.0009885	0.997
Alexander (1985)	3.063	-0.01165	0.997
KITRAL System	2.233	-0.01031	0.984

Estimating the rate of spread

In general, the elliptical propagation model can be summarized using the ROS aligned with the main axes of the ellipse. The head rate of spread (HROS) represents the fire speed in the wind direction (Figure 2.1, $HROS = r(0^\circ)$). Similarly, the back rate of spread (BROS) and the focal flank rate of spread (FFROS) are defined as $BROS = r(180^\circ)$ and $FFROS = r(90^\circ)$, respectively. It should be noted that in the literature, $FROS$ is reserved for the ‘‘Flank rate of spread’’, that represents the speed of the flank fires through the center of the ellipse, instead of its focus. ROS values for other ϕ angles can be computed directly by parameterizing the ellipse with respect to its focus or using approximated distribution schemes exploiting the values of the ROS in the main four axes.

Given the elliptical propagation model, these components satisfy a series of relationships (see Figure 2.4) summarized in multiple equations. As an example, the $BROS$ can be calculated as a function of $HROS$, LB , and the expected WS , using the following equation in the previously mentioned systems:

$$BROS(WS) = \frac{HROS(WS)}{HB(WS)} \quad (2.6)$$

where HB , known as the head-to-back ratio, is obtained from the LB as follows:

$$HB(WS) = \frac{LB(WS) + (LB(WS)^2 - 1)^{0.5}}{LB(WS) - (LB(WS)^2 - 1)^{0.5}} \quad (2.7)$$

where we explicitly indicate the dependency of $HROS$, $BROS$, LB , and HB to the wind speed WS . In the same line, we note that we can calculate the $FROS$ by:

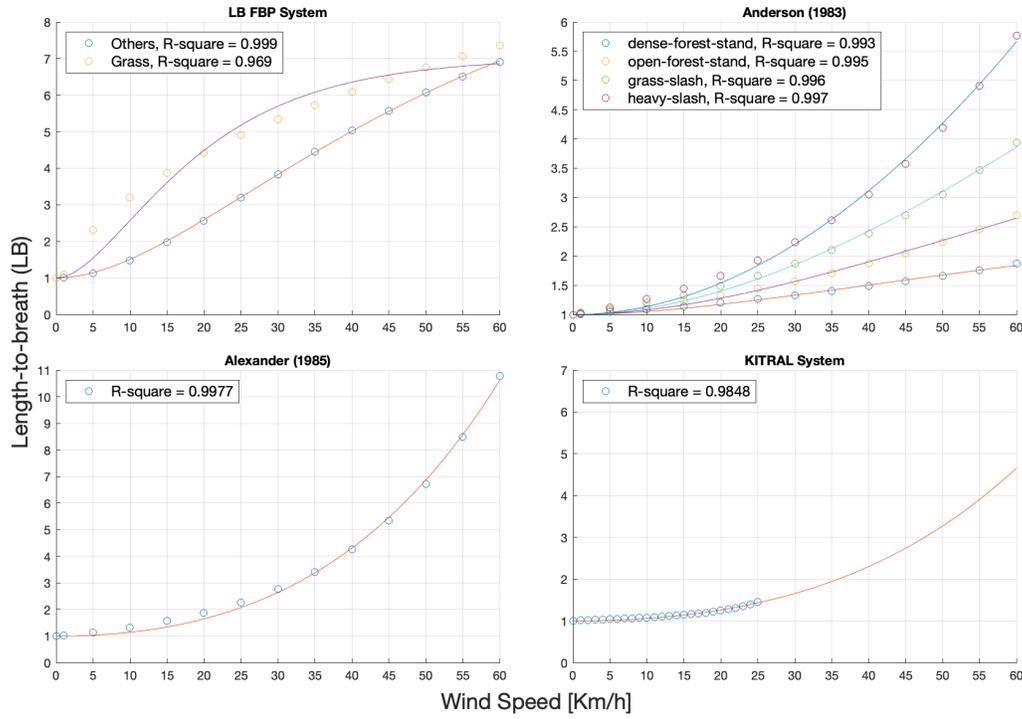


Figure 2.3: Fitted curves for different LB values (y-axis) as a function of the wind speed WS (x-axis) for the four fire spread models (Kitral, BEHAVE, FARSITE, and FBP). Data points are obtained from each fire spread model, given a set of predefined WS points to use during the curve fitting procedure. The optimal l_1 and l_2 parameters are estimated to maximize the coefficient of determination R^2 .

Then, $FROS$ is calculated by:

$$FROS(WS) = \frac{HROS(WS) + BROS(WS)}{2LB} \quad (2.8)$$

as can be clearly observed in the diagram presented in Figure 2.4.

Therefore, we can exploit the elliptical relationships of the propagation model to estimate some relevant components. As discussed by Van Wagner (1969) in [351] and Alexander (1985) in [14] the main elliptical components can be calculated in terms of the magnitudes defined

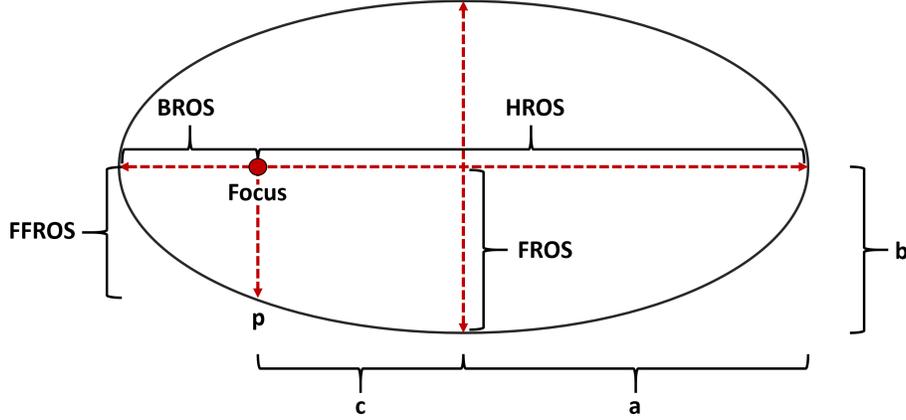


Figure 2.4: Elliptical propagation model components. The diagram illustrates the relationships between the elements of the model and the dimensions of the ellipse including the semi-major axis (a), the semi-minor axis (b), the distance from one focus to the center (c), and the semi-latus rectum (half the length of the chord through one focus, perpendicular to the major axis).

above and the elapsed time t since ignition, as follows:

$$a(W S, t) = \frac{HROS(W S) + BROS(W S)}{2} \cdot t \quad (2.9)$$

$$b(W S, t) = FROS(W S) \cdot t \quad (2.10)$$

$$c(W S, t) = \frac{HROS(W S) - BROS(W S)}{2} \cdot t \quad (2.11)$$

$$e(W S, t) = \frac{c(W S, t)}{a(W S, t)} \quad (2.12)$$

where a and b represent the length of the semi-major and semi-minor axes of the ellipse, respectively; c equals the distance from one focus to the center of the ellipse; and e is the eccentricity of the ellipse. We note that all these values are a function of the wind speed WS and the elapsed time since ignition, as the ellipse starts to expand/grow as a function of time t , simulating the dynamism of the fire propagation patterns.

With all these components, we have a comprehensive characterization of the ellipse representing the expected fire spread patterns at a surface level. Thus, we can easily obtain $ROS(\phi)$, the rate of spread for each angle ϕ , to use it in our fire growth model during the simulations.

Moisture scenarios

In practice, the effective ROS experience at a surface level mainly depends on the type of forest fuel that goes into combustion and the wind speed. However, another influencing factor is the moisture content of dead and live vegetation in the forest. This variable is crucial for estimating the surface ROS as certain land covers (e.g., grass types) tend to significantly modify their propagation patterns depending on their humidity levels. Thus, perturbing the dimensions of the estimated ellipse to calculate the ROS. With this in mind, multiple thresholds for the moisture values have been empirically studied and determined, to capture and simulate interesting situations for the practitioners and researchers. According to Scott & Burgan [307], interesting humidity levels can be represented by four main scenarios denoted D1L1, D2L2, D3L3, and D4L4. These are ordered from the driest (D1L1), where the fire tends to propagate faster, to the wettest (D4L4), with the opposite effect. This is translated into fires reaching higher $HROS$ values for similar WS as we move from wet to dry scenarios (Figure 2.5).

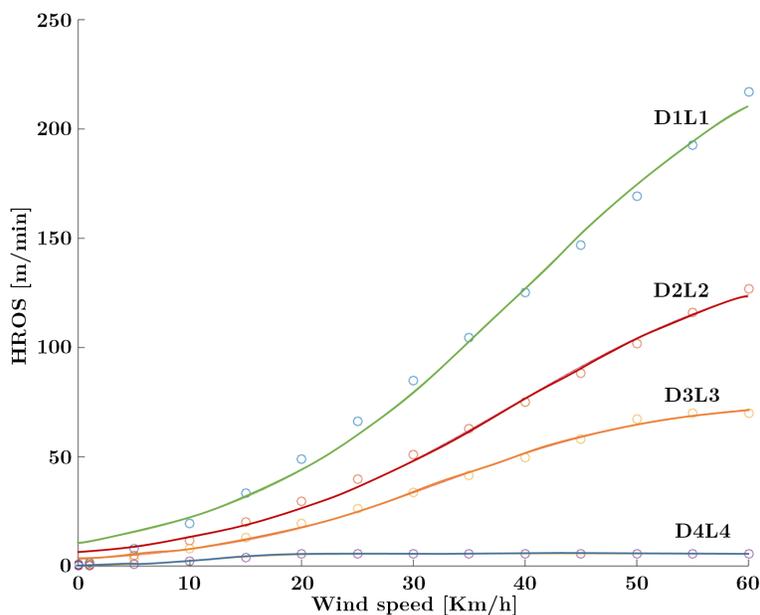


Figure 2.5: Example of fitted scenario curves for different $HROS$ values (y-axis) as a function of WS (x-axis). We use the a non-linear fitting method to adjust $HROS$ curves to the moisture scenarios. Data points for the fitting procedure are obtained from empirical results introduced in [307].

Based on the systems studied and the results obtained from the fitting procedure, the

following flexible head rate of spread (*HROS*) formula is proposed:

$$HROS(W S) = \frac{1}{p_1 \exp(-p_2 W S) + p_3} \tag{2.13}$$

where we estimate the *HROS* as a function of the wind speed $W S$ [km/h] and no slope effect is explicitly included. The $\vec{p} = (p_1, p_2, p_3)$ parameters are associated with each fuel type, i.e., each fuel type could have a different \vec{p} controlling the perturbation of the ellipse in the main wind direction to account for different propagation patterns depending on the land-cover. As an extension, a similar formulation replacing the \vec{p} vector by a $\vec{q} = (q_1, q_2, q_3)$ set of parameters could be used to model the effects of the crown fire phenomenon. In this case, the fire reaches the crown of the tree and starts to propagate towards adjacent trees, from crown to crown. This normally entails a faster rate of spread, leading to significantly more dangerous and severe wildfires. Studying and modeling this phenomenon requires training samples that recorded the exact moment when the fire reaches the crown to avoid bias. However, access to such information is complex and its quality varies significantly depending on the region of study, thus being out of the scope of this study.

Table 2.2: Example of *HROS* fitted parameters for all moisture scenarios, considering a land cover dominated by shrubs.

Scen. \ Param.	p_1	p_2	p_3	<i>R-square</i>
D1L1	0.08987	0.07819	0.003927	0.99
D2L2	0.1554	0.08058	0.006821	0.99
D3L3	0.3075	0.09702	0.01307	0.99
D4L4	4.902	0.2927	0.1805	0.99

Similar to the fitting procedure introduced in Section ??, we generate a series of data points representing the *HROS* for multiple $W S$ levels and scenarios using the existing fire spread models. Using the BehavePlus fire behavior system `andrews2014current`, which incorporates the definition of the four moisture scenarios and its expected *HROS* values for multiple wind speeds (0, 5, 10, 20, 30, 40, 50, and 60 km/h), we adjust the \vec{p} parameters introduced in eq. (13) using the NLS method described in Section ?? for each combination of land-cover and moisture scenario for all fuel types included in the FBP, Behave, and Kitral systems (a total of 316 independent fits).

Connecting all the previous modeling steps and structural components, we obtain a general fire spread model that i) can be adjusted by modifying its main parameters $\vec{l}b \in \mathbb{R}^2$ and $\vec{p} \in \mathbb{R}^3$ to reproduce multiple existing fire spread models, and ii) potentially trained to mimic propagation patterns observed in regions without existing spread models or previous studies. In this first work, we focus on the proposed methodology and development of the general fire spread model and tuning framework, leaving its large-scale training (with massive datasets of historical data) as a future project in our research.

2.3 Data Mining: Historical Fire Scars

In order to deal with cases where no regional publicly access datasets are available, we propose the following framework. Taking advantage of recent computational advances with global scale analysis capacity and the availability of planetary and historical satellite imagery in Google Earth Engine (GEE) catalog, we are able to generate historical fire scars from different regions. GEE is an open cloud-computing platform for geospatial analysis that contains a public catalog of satellite images, topography, land cover, and other environmental datasets [144]. GEE is an opportunity to generate a high-resolution database of wildfires. Taking advantage of the GEE big-data analysis platform, we develop a flexible workflow to reconstruct individual burned areas, topographic, and previous land cover data. This framework could be replicated by the international community with historical fire occurrence data or could be used in the reconstruction of recent global fires scars with free and public fire datasets including information about ignition coordinates, date, and duration, as in [24].

To obtain historical fire scars to train our spread model, we developed a semi-automated fire scar generator in GEE using a custom JavaScript: FireScar generator. This script automatizes the process of generating historical wildfire scars using the satellite (Landsat) imagery in the Google Earth Engine catalog. The workflow includes the following general steps/tools (Figure 2.6): i) input data selection and codification, ii) previous and post-fire satellite image mosaic generation, iii) preliminary burned area identification, iv) vectorization and spatial filtering, and v) output data generation of wildfire. It is important to note that the script can be used for any region in the world as long as its data are available.

The input data is a point pattern with geographical coordinates representing a (near) fire ignition point or a section within the burned area, indicating fire starting and extinction dates. Input data selection and codification are performed, gathering all the required data for the spatial location and temporal window of interest. The images searching engine includes two months of data before the fire starting date and two months after its extinction. We then compile both sets considering all images in that period and a 5 km. buffer around each ignition point.

Previous and post-fire satellite image mosaics are generated and processed by defining five functions to filter elements that degrade the quality of the image such snow, shadow, and clouds for each collection. In order to reduce burned area detection error caused by shadows, water bodies, agricultural or tree harvesting, we use the Normalized Burned Ratio (NBR, Eq. (2.14)) [178], the most common burned area index for burned area discrimination, and its multi-temporal form known as Delta Normalized Burn Ratio (dNBR) [88, 211]. Different burned area indexes are available. In this work, we use the Relative Delta Normalized Burn Ratio (RdNBR, Eq. (2.15)) because it has shown better results in our study areas [224].

$$NBR = \frac{(NIR - SWIR)}{(NIR + SWIR)} \quad (2.14)$$

$$RdNBR = \frac{(PreFireNBR - PostFireNBR)}{\sqrt{|PreFireNBR/1000|}} \quad (2.15)$$

where NIR is the near-infrared and the SWIR is the shortwave infrared wavelengths in Landsat satellite images. The preliminary burned area is identified, clipping the original mosaic.

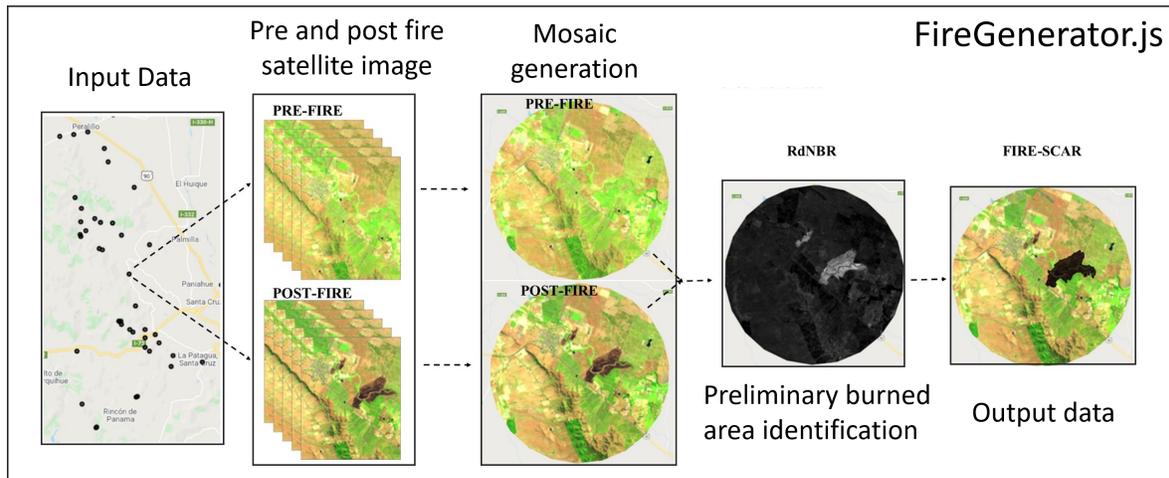


Figure 2.6: FireGenerator workflow. The main five steps of the script are depicted.

We select the best RdNBR index value for each wildfire that contains the burned area by visual interpretation. The raster mask in a binary form of the burned and not-burned pixel is converted to a vector format for spatial filtering. The individualization of the initially identified burned patches allows us to add spatial and spectral information to each one and discriminate between burned or available patches with new criteria, to diminish errors. The mosaic is then filtered by calculating the difference of pre and post indexes calculated from spectral bands, detecting those areas where the variation is higher, approximating the final fire scar

Layers of information are assigned to each individual burned patch: its size, the distance to the bigger patch, the distance to the ignition point, the maximum, minimum and mean value RdNBR, and the normalized difference vegetation index NDVI [211]. Finally, we export the fire scar perimeter for each wildfire including the land cover before the wildfire, elevation, slope, and most likely ignition area at a pixel level (30×30 m. resolution .tif files) and the weather conditions from the closes weather station available for the fire duration. This data is derived from the NASA Shuttle Radar Topographic Mission (SRTM), available in the GEE data catalog.

Additional data Sources

Land cover data for the Chilean instances were obtained from the most extensive and complete land covers dataset of Chile and publicly available by [372] in a 30 m. spatial resolution.

We obtain the hourly weather data from the nearest to wildfire meteorological station available from Red Agroclimatica Nacional (AGROMET¹). As input dataset, we used public databases of wildfires provided by multiple regional services (e.g., the Chilean Forest Service - CONAF²) that gives the spatial coordinates of ignition points and the ignition and fire control date. Similarly, input databases including ignition coordinates and fire spread dates were obtained from the Canadian Forest Service repository³, and directly from the Forest Science and Technology Centre of Catalunya⁴.

2.4 Learning Propagation Patterns

We formulate the learning propagation pattern model considering a vector $\vec{x} = (\vec{p}, \vec{l}\vec{b}, \vec{r})$ containing the adjustable parameters of the general fire spread model. The model is formulated for $|G|$ independent fuel types, with $\vec{p} \in \mathbb{R}^{3 \times |G|}$ the surface rate of spread factors, $\vec{l}\vec{b} \in \mathbb{R}^2$ the length-to-breadth ratio shape parameters of the elliptical model, and an extra set of adjustable weights $\vec{r} \in \mathbb{R}^3$ to fine-tune the elliptical ROS components ($HROS$, $BROS$, and $FFROS$) allowing the model to slightly deviate from elliptical propagation patterns, if needed. We define $t \in T$ the discrete set of time-steps from the ignition until burnt-out of a fire and $S_t(\vec{x}, F_z^i) : \mathbb{R}^{3 \times |G| + 5} \rightarrow \{0, 1\}^{m \times n}$ the simulator function that maps the vector of adjustable parameters \vec{x} applied to the i -th instance (landscape) of region z , F_z^i , into an $m \times n$ binary matrix (where 1 equals burned and 0 otherwise) representing the status of the landscape by time step t . We denote $\Pi_{t,z}^i$, the binary matrix associated with the real fire scar (training data) at time t from the $i \in N_z$ fire at region z .

For each region $z \in Z$ with N_z fires, we minimize the total fire growth evolution error defined as the difference between the set of simulated fires $\{S_t(\vec{x}, F_z^i)\}_{i=1}^{N_z}$ and the real scars $\{\Pi_{t,z}^i\}_{i=1}^{N_z}$, by obtaining \vec{x}^* , the vector of optimal parameters that minimizes the deviation between both scars. A set of initial values $\vec{x}_{0,z}$ are obtained from the elliptical adaptation of existing regional spread models, following the procedure described in Section ???. We formalize the optimization model as follows.

Definition 1 Let $\varepsilon_{z,T}^i(\vec{x}, F_z^i, \vec{\mu})$ be the error function for the i -th fire at zone z with duration T_z^i , $S_t(\vec{x}, F_z^i)$ the simulator function at time-step t , \vec{x} the adjustable parameters vector, and F_z^i the i -th forest instance from region z . Let $\Pi_{t,z}^i$ be the real-fire binary matrix (training instance) at period t , for the i -th fire of region z , $\|\cdot\|$ a matrix norm function (e.g., Frobenius norm), and $\vec{\mu}_i = (\mu_1, \dots, \mu_{T_z^i})$, an optional vector of weights associated with each time-step t of the fire propagation such that $\sum_{t=1}^{T_z^i} \mu_t = 1$. We define the z -Derivative-Free Optimization

¹<https://www.agromet.cl/datos-historicos>

²<http://www.conaf.cl/conaf/seccion-stadisticas-historicas.html>

³<http://FireGrowthModel.ca>

⁴<http://www.ctfc.cat/en/>

problem ($DFOP_z$) for an individual fire $i \in N_z$ as:

$$(DFOP_z) \quad \min_{\vec{x} \in \mathbb{R}^{3 \times |G| + 5}} \varepsilon_{z,T}^i(\vec{x}, F_z^i, \vec{\mu}_i) := \sum_{t=1}^{T_z^i} \mu_t \|S_t(\vec{x}, F_z^i) - \Pi_{t,z}^i\| \quad (2.16)$$

where $\|S_t(\vec{x}, F_z^i) - \Pi_t\|$ measures the error between the simulated and real scars at time-step t when using the parameters \vec{x} . Notice that $\sum_{r,c} s_{t_1,r,c}(\vec{x}) \leq \sum_{r,c} s_{t_2,r,c}(\vec{x})$ and $\sum_{r,c} \pi_{t_1,r,c} \leq \sum_{r,c} \pi_{t_2,r,c}$, if $t_1 \leq t_2$ since the burned area of a fire is proportional to the elapsed time. Weight vectors $\vec{\mu}_i$ are included to provide flexibility for different goals like minimizing the global hourly error or focus the adjustment on the final scar by setting $\mu_{T_z^i} = 1$. The optimal solution of $DFOP_z$ for the i -th fire of the region z , $\vec{x}_z^{*,i}$, provides the optimal parameters to reproduce the Π_z^i fire scar with minimum error.

To adjust the fire spread model to a region z , a set of N_z fires is used for training the main parameters of our model until convergence. Using the introduced $DFOP_z$ formulation, we define the Regional DFO problem at zone z ($RDFOP_z$) as follows.

Definition 2 Let $E_z(\vec{x}, \vec{\mu})$ be the z -regional error function with parameters \vec{x} , the vector of individual weight vectors $\vec{\mu} = \{\vec{\mu}_i\}_{i=1}^{|N_z|}$ to prioritize certain time-steps of the different regional fires, and $\vec{\omega} \in \mathbb{R}^{|N_z|}$ the vector of weights representing the individual contribution to the z -regional error by the i -th fire registered in the region, $i \in N_z$. Then, the $RDFOP_z$ is defined as:

$$(RDFOP_z) \quad \min_{\vec{x} \in \mathbb{R}^{3 \times |G| + 5}} E_z(\vec{x}, \vec{\mu}) := \sum_{i=1}^{|N_z|} \omega_z^i \varepsilon_{z,T}^i(\vec{x}, F_z^i, \vec{\mu}_i) \quad (2.17)$$

In $RDFOP_z$, we seek to minimize the total (additive) error of the individual fires in the region given a vector of fire spread parameters \vec{x} and the set of user-defined $\vec{\mu}$ and $\vec{\omega}$ weights. This last vector allows the fire scientist to provide specific weights to certain fires $i \in N_z$, in an attempt to capture the dynamics of – potentially – more relevant/useful events depending on the aims of the research study (e.g., large fires with crown fire, fires impacting native forest/population). After solving the $RDFOP_z$, we obtain the optimal set of regional parameters \vec{x}_z^* .

These optimization problem cannot be solved with conventional methods because we do not have information about the derivatives or the algebraic structure of the objective function. In order to solve this problem, we apply Derivative-Free optimization algorithms following the techniques and recommendations from [95] as discussed in Section 2.5.

Alternative error functions

In order to provide a flexible framework for researcher and practitioners, different error functions $\varepsilon(\cdot)$ are implemented to cover multiple objectives. Thus, researchers are provided with seven error functions including Frobenius norm (default), absolute difference between

scars, weighted norm, mean squared error, and structural similarity index [354]. In addition, we include extensions adding extra regularization terms (L_1 -norm and L_2 -norm) penalizing parameters with large values to obtain more stable fire spread models:

$$Norm = \|\Delta\| \quad (2.18)$$

$$W - Norm = \|\Delta\| / \sum_{r,c} \Pi_{r,c} \quad (2.19)$$

$$Absolute = SUM(|\Delta|) \quad (2.20)$$

$$W - Absolute = SUM(|\Delta|) / \sum_{r,c} \Pi_{r,c} \quad (2.21)$$

$$LFires = 0.5SUM^+(\Delta) - SUM^-(\Delta) \quad (2.22)$$

$$SFires = SUM^+(\Delta) - 0.5SUM^-(\Delta) \quad (2.23)$$

$$Burned = |SUM^+(\Pi) - SUM^+(S(\vec{x}))| \quad (2.24)$$

where $\Delta = S(\vec{x}, F) - \Pi$, the difference between the simulated and the historical fire scar (binary matrices); m and n the dimensions of the landscape, and $SUM()^{+,-}$ an auxiliary function to sum all elements of the matrix ($SUM(A) = \sum_{i,j} a_{i,j}$, $A \in \mathbb{R}^{I \times J}$), where $+$ and $-$ indicate that the summation is restricted to only positive or negative components, respectively. We also include the structural similarity index (SSIM) and mean square error (MSE) as alternative error functions.

As previously mentioned, a regularization term can be included in all previous functions. For this, we provide a λ parameter (penalty term) and the type of regularization (L_1 -norm or L_2 -norm) to be added in as part of the objective function:

$$\lambda \|(\vec{p}, \vec{q}, \vec{l}b)\|_L, \quad L \in L_1, L_2 \quad (2.25)$$

Users may design their own custom error functions by simply modifying the open-source code provided in the project's repository.

2.5 Derivative-free optimization

Derivative-free optimization (DFO) is an area of nonlinear optimization that deals with problems where the derivatives of the objective function (and potentially, constraints) are not available. Due to a growing number of applications in science and engineering, the development of DFO algorithms has increased and given greater attention in recent decades. Some applications using DFO algorithms can be found in [8, 16, 42, 153].

There are different situations where this methodology is appropriate: i) the functions defining the problem are provided through a computer simulation that cannot be easily subjected to automatic differentiation; ii) the optimization problem involves conducting a laboratory experiment, with no explicit mathematical expressions; iii) the objective function is noisy and the gradient estimation may be completely useless; iv) when the evaluation of

the functions require a significant amount of computational power, it may be prohibitive to perform the necessary number of function evaluations – normally no less than the number of variables plus one – to provide a single gradient estimation.

The diversity of applications includes problems in engineering, mathematics, physics, chemistry, economics, finance, medicine, transportation, computer science, business, and operations research (see e.g., [33, 95]). Some examples of them are: tuning of algorithmic parameters [34, 42]; engineering design [53, 54]; molecular geometry [9]; medical image registration [247]; and dynamic pricing [201].

In our research, we follow the ideas of parameter fit presented in [16] and [34]. Most numerical codes (for simulation, optimization, estimation, etc) depend on a number of internal parameters. Researchers implementing numerical algorithms know how critical the choices of these parameters are and how much they influence the performance of solvers. Typically, these parameters are set to values that either have some mathematical justification or satisfactory empirical results. One way to automate the choice of the parameters – in order to find possibly optimal values – is to consider an optimization problem whose variables are the parameters and whose objective function measures the performance of the solver for a given set of parameters, measured by CPU time or by some other indicator such as the number of iterations taken by the solver (see [95]). However, in our study, the main performance measurement is not CPU time or the number of iterations that Cell2Fire makes to get a more accurate fire scar. Since Cell2Fire simulated scars depend dynamically – in simulation time – on the fuel type (ROS obtained from the spread model), our parameters to adjust/re-scale the fire spread model will be such that they change the magnitude of the ROS among the main propagation axes. This way, our main performance measurement will be the adjustment error with respect to a real/historical fire scar observed, a real-time fire scar provided to predict the most likely evolution of an on-going fire, or a scar simulated by an already calibrated software.

In order to solve the optimization problem to adjust the parameters of the model, we apply a series of powerful and easy-to-implement DFO algorithms following the techniques and recommendations from [33] and [95]. Based on the characteristics of our problem and the expected performance of the different algorithms (convergence to the global optimum is not guaranteed), we implement, test, and compare the following algorithms in order to find the optimal parameters:

- **DFO:** This algorithm is based on approximating the objective function by a quadratic polynomial using predefined set points called “interpolation set” [94, 96]. Every n -dimensional quadratic model has $1 + n + n(n + 1)/2 = (n + 1)(n + 2)/2$ parameters, the DFO algorithm uses this number of interpolation points to determine them solving a system of linear equations. Then, assuming that the quadratic model approximates the objective function well within a certain region of a given radius, a new point is computed - within of a Trust-Region framework - to obtain a better objective function value. Later, a test measures how much has been achieved in reducing the objective function, compared to how much we reduced the quadratic model. Finally, an update of

the interpolation set occurs if necessary to improve the approximation of the objective function or a reduction of radius is effected.

- **Nelder-Mead**: an algorithm introduced in [243], it starts with a set of points that form a simplex – a generalization of the notion of a triangle or tetrahedron to arbitrary dimensions. On each iteration, the objective function values at the corner points of the simplex determine the worst corner point. The algorithm attempts to replace the worst point by introducing a new vertex in a way that results in a new simplex. Candidate replacement points are obtained by transforming the worst vertex through a number of operations around the centroid of the current simplex: reflection, expansion, inside, and outside contractions.
- **COBYLA**: this algorithm was developed to solve non-linearly constrained optimization problems [269]. This algorithm follows an approach similar to the DFO method [94, 96], but it uses a linear model approximation for the objective function and constraints, interpolating at the vertices that form a simplex and where a trust-region bound restricts the variables perturbation. Thus, a new vector of variables is calculated which may replace one of the current vertices, either to improve the shape of the simplex or because it is the best vector that has been found so far according to a merit function that gives attention to the greatest constraint violation. The trust-region radius is never increased, and it is reduced when the approximations of a well-conditioned simplex fail to yield improvement to the variables until the radius reaches a prescribed tolerance value that controls the final accuracy.
- **NEWUOA**: is an unconstrained optimization method using a quadratic interpolation approximation. Like the DFO method, it seeks to calculate the least value of an objective function by applying the trust-region iteration for adjusting the variables. Now, as we mentioned above, all n -dimensional quadratic models have $(n + 1)(n + 2)/2$ parameters. This means that, unless other conditions are imposed, we require this number of interpolation points to build them. However, in NEWUOA this is an input parameter denoted by m . In [271], the author proposed to use a quadratic model relying on fewer than $(n + 1)(n + 2)/2$ interpolation points. The remaining degrees of freedom in the interpolation are determined by minimizing the change to the Hessian of the surrogate model between two consecutive iterations. The latter is an advantage since a DFO algorithm aims to use fewer evaluations of the objective function.
- **BOBYQA**: is an iterative algorithm for finding a minimum of an $n - dimensional$ function subject to box-constraint. BOBYQA is a extension of NEWUOA, based on a quadratic interpolation approximation (see [270]).

The development of derivative-free algorithms dates back to the works of [243] and [317] with their simplex-based algorithms. An excellent review and numerical comparisons of state-of-the-art algorithms can be found in [95].

Table 2.3: Summary of the instances considered including total area in hectares, topographic characteristics (mean and range of the elevation in meters), fire duration in hours, dominant land-covers, and the total number of fuel types per instance.

Instance	Area [ha.]	Mean elev. [m]	Elev. range [m]	Duration [hr.]	Dominant Fuel	# Fuel types
CAN - Arrowhead	159,963	1329	[430,2894]	12	Red/White Pine	11
CAN - Central Kootenay	399,401	1811	[300,3000]	12	Boreal Spruce	11
CAN - Mica Creek	304,781	1523	[516,2895]	12	Red/White Pine	11
CAN - Revelstoke	350,956	1494	[429,2997]	12	Red/White Pine	11
CHI - Valparaiso	152	210	[14, 411]	27	Dense shrubs	8
CHI - Angol	182	156	[62,600]	71	Old Conifers	8
CHI - Carahue	337	439	[63,765]	54	Native forest	7
CHI - Lota	182	97	[0,463]	21	Dense shrubs	8
SPA - NautAran	219	1640	[969, 2532]	13	Grass and shrubs	27
SPA - Poblamontornes	186	194	[58, 334]	20	Grass and shrubs	22
SPA - Valbona	212	537	[0,720]	13	Grass and shrubs	24
SPA - Jonquera	1530	187	[0,947]	117	Woody shrubs	32

2.6 Experiments

Data Description

Historical scars. Three main set of instances consisting of four fires from different zones including forest land cover, topography, and hourly wind stream files are tested: (i) Canadian forests from the British Columbia region, (ii) Spanish landscapes located in Catalunya, and (iii) South of Chile. These fires were collected and processed from the Canadian Forest Service repository⁵, the Forest Science and Technology Centre of Catalunya (Spain), the National Forest Corporation (Chile), the public Fire Atlas dataset (2019) [24], and using our proposed data mining framework.

Homogeneous forests. Two sets of 10,000 ha. homogeneous landscapes covering all fuel types available in state-of-the-art simulators for Canada (Prometheus, 18 types) and Spain (Farsite, 32 types) are generated – a total of 50 forest. This set is used to show how our proposed general spread model is able to mimic the propagation patterns of currently validated simulation models by individually tuning our model for each fuel type, being able to obtain comparable results to these systems with a significant computational performance advantage.

All data is available in the public repository of this project⁶. Historic instances were affected by suppression activities (e.g., firefighters interventions). A detailed description of all instances can be found in Table 2.3.

⁵<http://FireGrowthModel.ca>

⁶<https://github.com/cpaismz89/GlobalCell2Fire>

Algorithms

Following the previous Chapter, we integrate our Cell2Fire fire-growth simulator (see Chapter 1) with the derivative-free optimization automatic adjustment algorithms via an optimization module that aims to minimize the differences between the simulated wildfires and the expected ones, solving a series of optimization problems and processing all relevant inputs/outputs.

We implement the derivative-free algorithms described in Section 2.5 directly in Python using the NLOPT [167] package to solve the optimization problems. Best results were obtained using the following strategy: (0) p , q , and lb initial values are set using the fitting scheme described in Section 2. (1) Starting points $\vec{x}_{0,z}$ for the optimization problems are obtained after a simple grid-search procedure to minimize the global error ($RDFOP_z$) modifying the first (head ROS) and last (eccentricity) factors of the elliptical dimensions r inside the $[0, 3]$ interval using a step of 0.2 units. (2) After selecting the starting point $\vec{x}_{0,z}$, we first solve the $RDFOP_z$ problem for r (with fixed p , q , and lb) obtaining $\vec{x}_{0,z}$. (3) Then, we solve the original $RDFOP_z$ problem allowing the algorithms to modify any parameter. (4) As ending criteria, we set an absolute tolerance on the variation of the current solution at iteration n , $|\Delta x_n| < xtol_abs = 1e^{-26}$, a maximum number of $N = 5000$ evaluations, or 5 hours of running time. If supported by the algorithm (e.g., BOBYQA), the following box constraints are provided for each set of parameters: $r \in [0, 10]$, $p, q \in [-1, 12]$, and $lb \in [-1, 5]$.

All experiments have been conducted in a laptop with a 4th generation I7 CPU (1.9 GHz, 2 cores), 8 GB of RAM, and Ubuntu 14.0 OS.

Evaluation

We focus on a tactical landscape planning perspective (prevention models). Therefore, we minimize the z -regional error with respect to the final scars of the N_z fires, not accounting for deviations in their temporal evolution. We set $\mu_t = 0 \forall t < T_z$ and $\mu_{T_z} = 1, \forall z \in Z$ obtaining $\varepsilon_{z,T}^i(\vec{x}, F_z^i) = \|S_T(\vec{x}, F_z^i) - II_{T,z}^i\| \forall i \in N_z$. We set $\omega_z^i = 1/A_z^i \forall i \in N_z$ with $A_z^i := \sum_{r,c} \pi_{T,z,(r,c)}^i$, the total number of ones inside the final scar binary matrix to normalize the contribution of each individual error function $\varepsilon_{z,T}^i$, accounting for different dimensions of the fire scar and land. This setting allows us to avoid bias when optimizing \vec{x} , penalizing the objective function proportionally to the error incurred on each instance with respect to the historical fire scar.

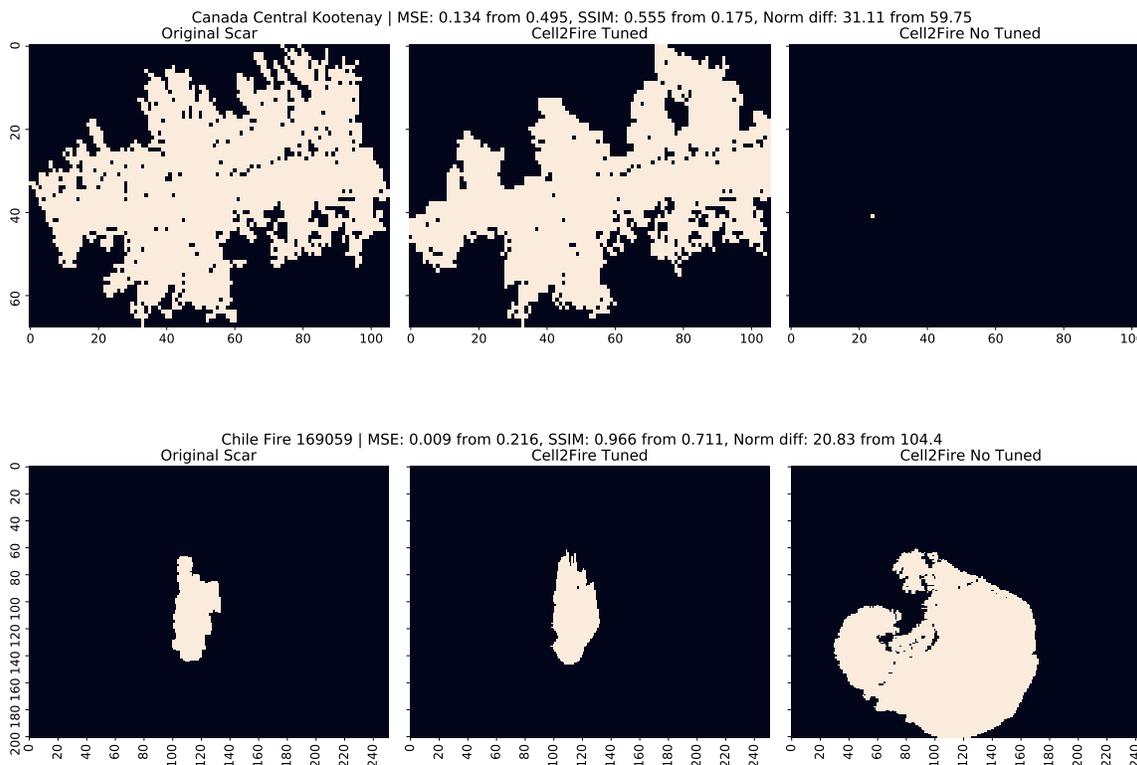
We measure the performance of the learning framework by reporting the optimal value of $RDFOP_z$ and providing an individual analysis of the N_z fires indicating the best mean squared error (MSE), structural similarity index (SSIM), and the norm of the difference between the simulated and real scars $\delta_{Norm_z}^i = \|II_{T,z}^i - S_T(\vec{x}_z^{*,i}, F_z^i)\|$. Best results for all instances converging to optimal \vec{x}^* values starting from 100 random variations of the original point \vec{x}_0 using a normal perturbation $\xi \sim N(0, 4)$ are reported.

2.7 Fire spread: learning & adjustments

Based on the experiments, best results in convergence and number of evaluations during training are obtained with the BOBYQA algorithm using the configuration described in the previous section. Thus, we focus our discussion on the results obtained from its application.

Historical Scars

As discussed in the Introduction, we aim to provide the research community with an effective and efficient fire spread learning framework, flexible enough to mimic the general propagation patterns observed in real life to obtain relevant insights for tactical decision-making process such as fuel-treatment plans, resource allocation priorities, among other potential actions to mitigate the expected losses due to wildfire risk.



Reproducing historical wildfires is a challenging task. Multiple approximations of relevant inputs such as land cover classification and weather streams as well as lack of information about the presence of suppression activities during the evolution of the event introduce several sources of error that play against the accuracy of simulators. Besides these major challenges, our adjusted fire spread model is able to accurately reproduce historical wildfire scars in all the three regions tested. From the experiments, the adjustment to Chilean wildfires tends to be faster (less iterations and evaluations), more accurate, and starting from

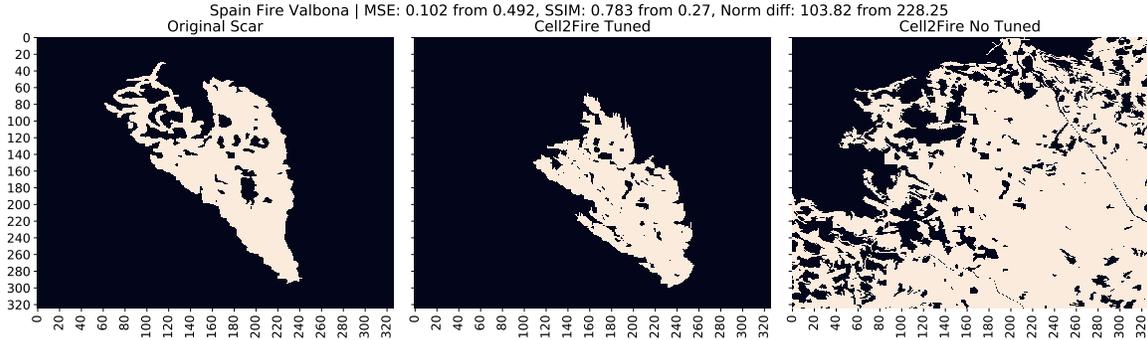


Figure 2.7: Samples of Canada, Chile, and Spain original fire scars (left), simulated fire scars post learning (\vec{x}^* - center), and simulated fire scars using default parameters (\vec{x}_0 - right). Individual MSE, SSIM, and δ_{norm} metrics are reported.

high-quality points as shown in Table 2.4. This is mainly explained by the following characteristics: (1) Chilean instances tend to present less to none external suppression interventions since the selected fires were mostly naturally suppressed and (2) the number of different fuel types is less than in the Canadian and Spanish cases, introducing fewer discontinuities in the landscape and thus, less noise when solving the optimization models. Analyzing the simulated wildfires of the Canadian and Spanish landscapes, we observe significant improvements in both regions, reaching differences up to 76% and 79% (MSE), 88% and 66% (SSIM), and 59% and 55% (δ_{norm}) between the final adjusted scar (\vec{x}^*) and the starting point of the optimization (\vec{x}_0).

With average SSIM values of 0.64 (Canada), 0.95 (Chile), and 0.73 (Spain), our adjustable fire spread model is able to capture the main propagation patterns of three different regions, with only a limited dataset of historic wildfires (Figure 2.7, Table 2.4). All adjusted simulated scars can be found in Figure 2.8.

Homogeneous Forests

From the homogeneous experiments (Tables 2.5 and 2.6), we are able to obtain similar propagation patterns to the ones observed in both Prometheus (Canada) and Farsite (Spain) for all independent fuel types, allowing us to mimic the performance of these state-of-the-art simulators by simply adjusting the r parameters of our model and keeping the starting p , q , and lb values obtained from the fitting procedure fixed. From the results, final average SSIM of 87% and 93% are obtained for both regions, converging to the optimal solution in less than one hour of training (an average of 89 evaluations). This independent tuning framework allows us to obtain high-quality parameters and starting points to accurately reproduce simulations obtained by these two validated systems as well as showing the adaptability and

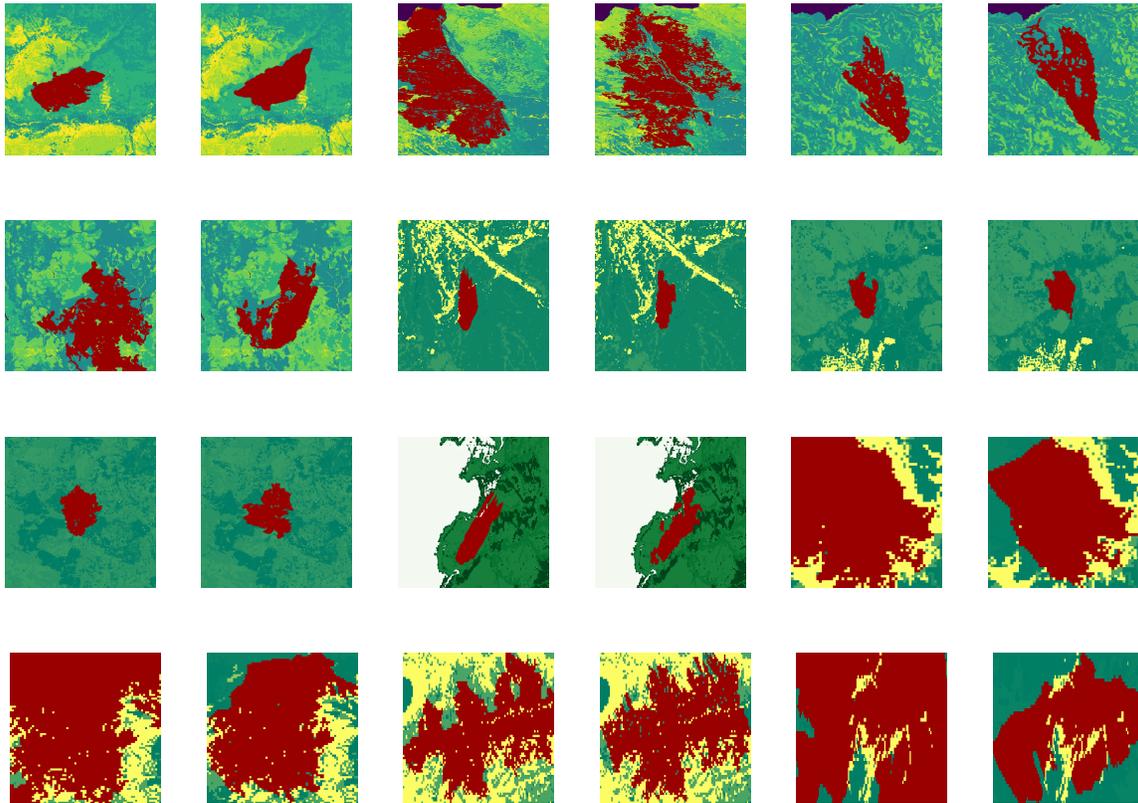


Figure 2.8: Final real and simulated wildfire scars comparison. Historical wildfires for the different instances are shown in even columns and simulated scars after applying the adjustment framework to the general fire spread model are depicted in odd columns.

flexibility of our spread model to be adjusted to different propagation patterns and forest compositions.

This way, researchers familiarized with the propagation patterns obtained with existing simulation models would be able to reproduce their projects and use our spread model as a starting point for further improvements as well as taking advantage of all the outputs and statistical tools provided with our system. Moreover, researchers from regions where no existing/validated spread models are available would be able to develop their own custom spread models by providing empirical observations or lab experiments of propagation patterns for multiple fuel types as the main training data for our model, learning new propagation patterns for currently unsupported land covers by existing simulators, providing them with access to high-quality simulation tools.

Table 2.4: Best results obtained for each instance. Optimal values for MSE , $SSIM$, and δ_{norm} are indicated (*) and compared with their initial values (0) using \vec{x}_0 as starting point.

Instance	MSE^*	MSE_0	$SSIM^*$	$SSIM_0$	δ_{norm}^*	δ_{norm}^0
CAN - Arrow Head	0.163	0.549	0.698	0.118	30.68	56.31
CAN - Central Kootenay	0.134	0.495	0.555	0.175	31.11	59.75
CAN - Mica Creek	0.125	0.512	0.788	0.091	19.80	40.05
CAN - Revelstoke	0.325	0.489	0.418	0.245	43.20	53.00
CHI - Valparaiso	0.009	0.216	0.966	0.711	20.83	104.40
CHI - Angol	0.013	0.044	0.963	0.935	28.57	51.69
CHI - Carahue	0.032	0.065	0.925	0.724	60.37	151.60
CHI - Lota	0.025	0.065	0.930	0.895	38.70	62.63
SPA - NautAran	0.055	0.086	0.905	0.854	77.76	96.87
SPA - Poblamontornes	0.162	0.386	0.714	0.454	122.68	189.16
SPA - Valbona	0.102	0.492	0.783	0.27	103.82	228.25
SPA - Jonquera	0.305	0.383	0.510	0.389	482.66	541.19

Sensitivity

Given the number of adjustable parameters present in our general spread model, multiple combinations can lead to similar spread patterns. To measure the quality and robustness of the obtained solutions, we analyze the impact of perturbing each converged parameter (individually) by a random noise $\eta \sim N(0, 2)$, keeping those solutions where the variation in performance (MSE, SSIM, and δ_{norm}) is less than 5% in any metric. This way, the final spread model will be more flexible and able to better generalize to new observations. Researchers are encouraged to perform similar analyses when training the spread model, avoiding unstable solutions that can lead to inaccurate simulations and potentially ineffective decisions. To this respect, the experimentation and application of error functions including regularization terms is suggested to avoid fire spread models that depend on a small subset of parameters.

Based on our experiments, solutions tend to converge faster when training the general spread model in a two-step approach: (1) Fixing initial p and lb values based on regional spread models or selecting the closest one to the region of interest while adjusting the r vector and (2) fine-tune these values after fixing the r vector. This allows the optimization method to better capture the impact of each parameter in the objective function, improving and approximating the trust-region for the most relevant parameters as well as providing useful bounds for the researcher when calling the DFO algorithms such as BOBYQA, reducing the feasible region and improving their convergence. Furthermore, this becomes crucial when dealing with fragmented landscapes including multiple fuel types. In those cases, convergence can be slow and inaccurate due to the number of variables in the model. Based on preliminary

experiments, we obtained the best performance after (1) detecting the most frequent land covers under the fire scar, (2) adjusting their associated parameters while keeping the rest fixed, and (3) adjust the remaining factors. Several tuning schemes can be tested, a topic that remains an open challenge for future research.

2.8 Conclusions

In this Chapter, we have developed an effective and flexible learning methodology for training a general fire spread model with historical data. Starting from simple adaptations of regional fire spread models, we were able to simulate realistic fire scars of three different regions of the world. To the best of the authors' knowledge, this work represents the first automatized open-source learning system in the context of global wildfire simulation.

Taking advantage of the massive global datasets and increasing computational resources, our framework is an opportunity for countries lacking fire behavior models. As we discussed, the development of fire behavior models is challenging and could take decades and high economic cost. This Chapter proposes a framework for fire behavior model fitting that needs only four inputs i) ignition point, ii) land cover data, iii) local topography data, and iv) atmospheric conditions during the wildfire. Each of these data could be obtained freely from different data servers and complemented with regional datasets. For example, to deal with cases where no local publicly accessible data sets are available, GEE is a valuable platform for accessing geo-databases and prepare them for the model. However, users must be aware of these global datasets attributes as well as the spatial, temporal, and thematic resolution to use the best data for model fitting purposes.

Our proposed framework could be replicated by the international community with national historical fire occurrence data, with recently available global data, or a combination of both. For example, in [24], the authors make freely available a global data set of individual fire scars with information of the size, duration, speed, and direction of the fire and estimate the ignition points location. A total of 13.3 million individual fire information between 2003–2016 at 500 m spatial resolution and a minimum wildfire area of 25 ha. On the other hand, [211] estimates individual fire-scars globally with a spatial resolution of 30 m using Landsat satellite image with promising results. Therefore, it is expected that massive data sets of high-resolution historical fire scars could be available for the entire globe shortly. In the same line, land cover data is incrementally available for the whole planet, including underdeveloped countries. The authors in [68] deliver a global land cover map at 100 m spatial resolution that provides an annual land cover for the 2015-2019 period with 22 identified land cover classes, including 16 different vegetation types. If required, more detailed land cover data could be found in country-level initiatives. In the absence of observed local conditions for meteorological data, ERA5-Land (fifth generation of the European Centre for Medium-Range Weather Forecasts-ECMWF atmospheric reanalyses of the global climate) combines physical models with field observation to estimate hourly atmospheric conditions from January 1981 to three months from real-time [294].

Therefore, our framework and available data sets are great opportunities for fire planning in underdeveloped countries and test available models with observed fire behavior in the absence of local fire growth models or economic resources to fit field-based behavior models.

The inclusion of explicit factors/thresholds to represent the effect of suppression activities remains an open challenge for future improvements. Given historical fire scars with explicit human interventions, the system could be able to determine thresholds associated with the ROS and fire intensity levels where suppression activities (e.g., performed by firefighters) could be applied based on the learning from past actions, detecting critical sections of the landscape and providing relevant information useful for decisions such as resource allocation. Future research would be focused on (a) extending the analysis to other regions, (b) develop an automatized online-learning system that updates the current optimal parameters when new events are registered, (c) characterize the behavior and relevant patterns behind the crown fire phenomenon – an open challenge for fire scientist – using historical data, and (d) integrate it as part of a tactical landscape management decision support system.

In the following Chapters, we will focus our attention in the integration of our wildfire simulation and adjustment models with complex data-driven decision making models for landscape planning, highlighting the powerful insights obtained by the interaction between all the components of our framework.

Table 2.5: Summary of Canadian fuel types adjustment.

Fuel	MSE^*	MSE_0	$SSIM^*$	$SSIM_0$	δ_{norm}^*	$\delta_{norm,0}$
C1	0.07	0.14	0.80	0.73	25.65	36.85
C2	0.09	0.10	0.77	0.77	29.21	31.50
C3	0.09	0.12	0.77	0.74	29.60	34.25
C4	0.09	0.34	0.77	0.55	29.15	58.57
C5	0.05	0.10	0.84	0.81	21.68	32.12
C6	0.05	0.21	0.85	0.70	21.73	45.31
C7	0.04	0.10	0.85	0.83	20.49	30.95
D1	0.01	0.02	0.99	0.92	5.34	14.49
D2	0.00	0.00	1.00	0.99	0.00	1.00
M1	0.02	0.21	0.89	0.68	26.45	45.71
M2	0.02	0.20	0.91	0.71	27.66	45.12
M3	0.04	0.35	0.88	0.52	26.55	58.87
M4	0.01	0.42	0.88	0.61	21.42	48.91
O1a	0.01	0.06	0.91	0.84	12.31	24.62
O1b	0.02	0.09	0.88	0.80	11.22	29.39
S1	0.03	0.09	0.85	0.79	13.51	29.41
S2	0.03	0.07	0.89	0.81	12.98	25.94
S3	0.01	0.09	0.89	0.77	10.44	30.74
Mean	0.04	0.15	0.87	0.75	19.19	34.65
STD	0.03	0.12	0.06	0.12	8.96	14.41

Table 2.6: Summary of Spanish fuel types adjustment.

Fuel	MSE^*	MSE_0	$SSIM^*$	$SSIM_0$	δ_{norm}^*	$\delta_{norm,0}$
101	0.02	0.07	0.86	0.72	5.39	10.63
102	0.01	0.16	0.93	0.77	4.80	15.97
103	0.01	0.06	0.96	0.90	4.00	9.80
104	0.00	0.00	1.00	1.00	0.00	0.00
105	0.00	0.00	1.00	1.00	0.00	0.00
106	0.00	0.00	1.00	1.00	0.00	0.00
107	0.00	0.00	1.00	1.00	0.00	0.00
108	0.00	0.00	1.00	1.00	0.00	0.00
121	0.02	0.52	0.84	0.30	6.08	28.88
122	0.03	0.32	0.91	0.58	6.63	22.67
123	0.01	0.11	0.96	0.83	4.47	13.45
124	0.01	0.22	0.93	0.71	4.58	18.57
142	0.01	0.41	0.93	0.16	3.46	25.67
143	0.01	0.17	0.93	0.53	3.61	16.25
144	0.01	0.17	0.95	0.77	3.87	16.46
145	0.01	0.05	0.96	0.91	3.74	9.17
146	0.06	0.26	0.71	0.53	9.64	20.57
147	0.01	0.17	0.96	0.76	3.61	16.61
148	0.04	0.24	0.88	0.57	7.87	19.61
149	0.01	0.11	0.97	0.84	3.74	13.15
161	0.00	0.14	0.97	0.59	2.45	14.76
162	0.02	0.62	0.86	0.09	6.00	31.58
163	0.03	0.24	0.91	0.66	6.40	19.70
164	0.01	0.76	0.92	0.02	3.87	34.91
165	0.02	0.30	0.86	0.03	5.83	21.98
181	0.00	0.00	0.96	0.96	1.41	1.41
182	0.00	0.00	1.00	0.93	0.00	2.45
183	0.00	0.01	1.00	0.92	0.00	3.32
185	0.00	0.22	0.96	0.47	2.24	18.79
186	0.01	0.40	0.94	0.21	3.16	25.20
188	0.01	0.30	0.95	0.36	3.00	21.77
189	0.01	0.44	0.93	0.14	4.00	26.40
Mean	0.01	0.20	0.93	0.63	3.56	14.99
STD	0.01	0.19	0.06	0.32	2.51	10.10

Chapter 3

Decision making under wildfire uncertainty: A data-driven optimization approach for landscape management

3.1 Introduction

In the previous two Chapters, we delved into data-driven approaches for wildfire simulation and adjustment of fire spread models using historical data to accurately represent potential wildfires in any region. As we have mentioned, the destructive potential of wildfires has been exacerbated by climate change, causing their frequencies and intensities to continuously increase globally. Therefore, generating fire-resilient landscapes via efficient and calculated landscape management plans is critical to protecting native forests, agricultural resources, biodiversity, and human communities. To tackle this challenge, in this Chapter, we propose a framework that integrates our previous models with optimization models. We introduce the concept of *Downstream Protection Value* (DPV), a flexible metric that assays and ranks the impact of treating a unit of the landscape, by modeling a forest as a network and the fire propagation as a tree graph. Using our open-source decision support system, custom performance metrics can be optimized to minimize wildfire losses, obtaining effective treatment plans focusing the treatment on the most critical zones.

The link between climate change and increased wildfire risk highlights the need to change the paradigm of how we coexist with fire and our environment. On the globe, it has been recorded that forest fires can be produced by lightning, volcanic activity, a spark from a rockfall, or by human carelessness [305]. For example, in Canada, lightning, human activity, and unknown causes account for 47%, 49%, and 4%, respectively, of all forest fires that occurred between 1990 and 2016, as presented in [340]. In the same line, the authors also point out that projected increases in burned areas suggest that the current state of forest fire management will not be able to adequately cope with this rate and warns that a paradigm shift in landscape management based on fire risk should occur. Research indicates that

these fire events will likely increase due to the effects of climate change on temperature, precipitation levels, and soil moisture, having increased the number of fires and area they consume around the globe, generating larger events than before [290, 359, 360].

The current incidents highlight the need for preemptive policy measures to reduce the risk of fire occurrence [102], managing the land in an effective way [73] to protect natural forests, agricultural areas, and human lives. These concepts are included in what is known as FireSmart Forest Management (FSFM) (see also [340]). As discussed in [158], this paradigm considers opportunities in three dimensions: i) decrease of the fire behavior potential of the landscape, ii) reduction of the potential for fire ignitions, and iii) increase in the fire suppression capability. However, the success of these challenges cannot be addressed individually, and coordination by actors, whether private, public or government is required as mentioned in the National Cohesive Wildland Fire Management long-term goals and strategy [99].

Several researchers have reported that the intensity and severity of wildfires can be reduced through fuel management activities. Targeted fuel treatments and firebreaks are used to alter the composition of wildland vegetation and forest in order to modify the behavior of future fires. There are multiple fuel treatment activities to delay the spread of a fire such as mechanical treatments, prescribed burns, thinning, and so forth [3]. These actions are mainly used to decrease the intensity and size of potential wildfires, but they could also affect species diversity [312], restore fire-dependent ecosystems, and help suppress future fires, among many others [116, 157, 230]. Research by [71, 112, 115, 116], indicates that it is possible to modify fire behavior and progress across landscapes through strategic placement of treatments. In [179], the authors develop and evaluate a series of management activities applied in certain patterns (dispersed, clumped, regular, and random patterns). They study the impact of these configurations on future wildfire behavior – flame length and fireline intensity – and their cumulative effect over multiple planning periods. The authors conclude, after analyzing the output of multiple simulations over decades for a case study located in Oregon, that the proposed approach can only marginally alter the size and severity of future wildfires under tight budgetary conditions. Other researchers [79, 129] have studied the benefit of placing a series of parallel strips acting as fuel breaks in the landscape to reduce the average fire spread rate when the fire propagates perpendicular to their placement. A model based on percolation theory [45] is introduced to evaluate the impact of spatially correlated (connected) treatments starting from random treatments to support suppression forces by allocating fuel breaks between wildland and developed areas. In [296], the authors discuss strategies to define treatments using burning probability (BP) maps, the fire size or the flame length; an application of this metric can be seen in [319], where burn probability; wildfire risk to restored seismic line areas; and the effectiveness of mitigation measures were calculated using the Burn-P3 program [257] (see an application of burn probability modeling review in [259]). Other studies have focused on finding the optimal spatial allocation of prescribed burning activities [11, 218] and the design of firebreaks to control fire spreading [292].

A pioneering work using a simulation-optimization approach allowing the authors to measure the impact of a landscape optimization heuristic under wildfire uncertainty and forest growth is found in [310]. In the same line, the authors in [293] design a computationally

intensive decision support system using a stochastic simulation-optimization approach to evaluate the performance of different fire-breaks allocations generated by a metaheuristic. Applying their framework to a case study in western Ontario, the authors identify useful patterns to minimize the expected fire-risk within the landscape, comparing their solutions with and outperforming simpler methods. An integrated fuel management and suppression planning model is presented in [225], where the authors propose an integrated integer optimization model incorporating both decisions. The authors show the advantages of the integrated planning approach to maximize the coverage of the landscape using a series of generated instances, providing useful insights for decision makers to balance suppression and prevention efforts. In [171], the authors present a two-stage stochastic integer programming method to select the optimal areas to perform fuel treatment, generating fire behavior data with FARSITE simulator [117]. This approach aims to minimize the total treatment and expected future losses/costs under weather uncertainty and forest growth in a study area located in East Texas. The authors analyze the trade-off between the elements of the objective function, showing that performing no fuel treatment can be the optimal decision depending on the costs involved.

Novel decision support frameworks to locate optimal treatment and examine trade-offs among alternative restoration strategies are proposed in [349] and [4]. Both studies consider objectives such that provisional ecosystem services, fire protection, and some key ecological stressors are included in the analysis using FlamMap wildfire simulation program finney2006overview to estimate the impact of the fire.

In [263], the authors report that the systematic segmentation of fuel can influence fire size and fire growth. However, the empirical evaluation of the efficiency of a firebreak and/or fuel treatment is difficult, if not impossible. Examples of some practical applications can be seen in [6, 332]. The weather conditions, as well as the state of a forest (e.g., surface fuel load, foliar moisture content) at any given moment, are non-reproducible. This unreliability in physical conditions is one of the advantages of deterministic fire simulators such as FARSITE, Prometheus and Wildfire Analyst [112, 273, 339], which can reproduce fires with and without treatment activities. However, these simulators do not have a well-structured interface to facilitate data exchange during multi-period interactions – e.g., to systematically evaluate fuel management in operational or tactical planning – and do not consider temporal changes of vegetable fuels. Additionally, evidence of the high stochasticity of this problem and inconsistency of forestry tools are highlighted in [39] where the authors show that, historically, the majority of treatments rarely intersect with wildfire occurrences. Therefore, finding an effective treatment layout across a landscape – even for a single time period – is not a simple task, and could be amplified according to the socio-ecological context in which fire management decisions are made [4, 331].

Recent work in [87] discusses challenges such as timing, selection of an adequate fuel-treatment method, and the impact of high uncertainty levels over multiple time periods (e.g., climate and ignition). There is evidence that prioritizing areas to treat based on non-fire risk metrics has not been effective [40]. Then, to fully address the optimal treatment problem, the development and systematic evaluation of effective fire risk metrics in an integrated,

robust, and flexible decision support system are required. In this research, we try to fill that gap, introducing different fire risk metrics to address the prioritization of firebreaks, including Burn Probability (BP), Betweenness Centrality (BC) [64], Fire Protection Value (FPV) [253], and a new metric proposed called Downstream Protection Value (DPV). DPV is a flexible metric that measures and ranks the impact of treating a unit of the landscape to minimize wildfire losses, similar to the “node influence grid (NIG)” introduced in FlamMap [114] but with substantial conceptual differences as will be discussed. Our objective is to select the treatment units or stands, defined as homogeneous sections of the forest with similar fuel type, age, and topographic conditions, with the highest impact on interrupting the fire propagation by incorporating DPV into an open-source, integrated Decision Support System (DSS), providing researchers with optimal landscape treatment plans under wildfire uncertainty. To assess the performance of our new methodology, we generate and apply the fuel treatment plans in two Canadian regions. We discuss the sensitivity of alternative metrics and DPV with respect to their learning capacity and performance under multiple sources of uncertainty (weather conditions and ignition). Their performance is measured by estimating (1) expected wildfire losses, (2) average fire propagation rate, and (3) their adaptability. A description of our proposed data-driven framework can be seen in Figure 3.1.

3.2 Methods

Study area

In this work, we use public data of real forest patches from Alberta and British Columbia provinces, Canada (Table 3.1). The Alberta instance is a sample data set of the Rocky Mountain subregion for the 2001 Dogrib fire¹. This patch has a surface of 79,611 ha, divided into 100×100 [m^2] cells. British Columbia instance is a very large section (1,854,838 ha) and it is characterized by mountainous surfaces and discontinuous fuel levels. In order to analyze the possible effects of topography on risk measures and accelerate our computational capacity, we build 6 sub-instances from it. We denoted them as ArrowHead (AH - 159,963 ha.), Revelstoke (RT - 350,956 ha.), Mica Creek (MC - 304,781 ha.), Glacier Natural Park (GNP - 464,664 ha.), Central Kootenay (CK - 399,401 ha.), and Neptune Peak (NP - 350,956 ha.). We picked this naming convention following the names of the mountains within each patch. The full British Columbia instance is part of the Burn-P3 documentation². Maps depicting all instances with their fuel models and topographic characteristics can be seen in Figure 3.2. Each instance contains all data layers (fuel composition, topographic characteristics, and weather conditions) for each cell.

¹http://www.firegrowthmodel.ca/prometheus/software_e.php

²http://www.firegrowthmodel.ca/burnp3/software_e.php

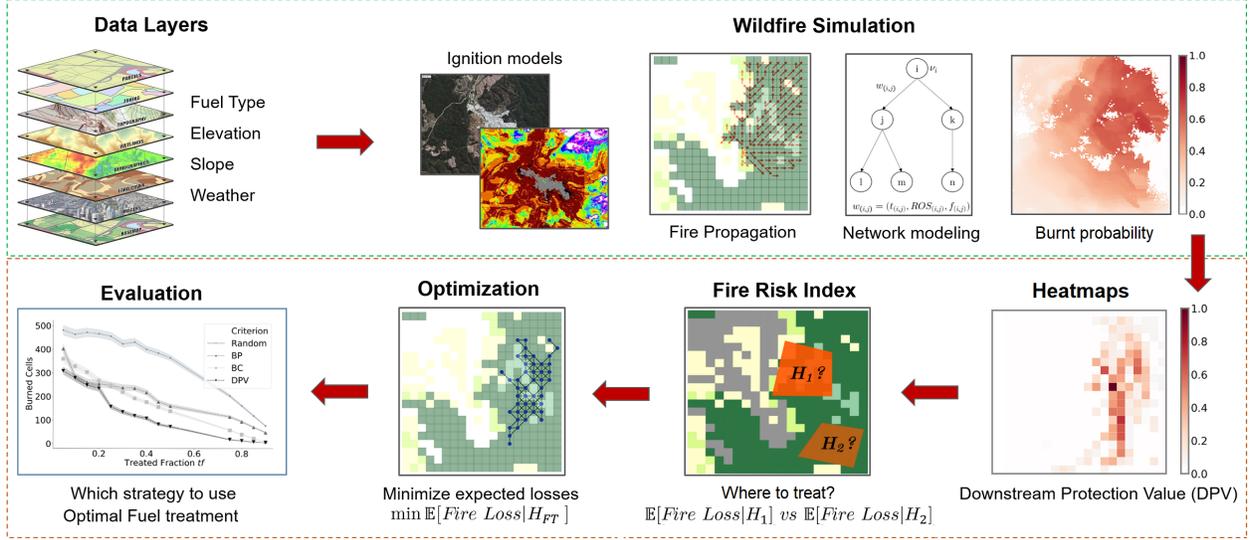


Figure 3.1: Framework schematic. (PHASE I) Two modules are integrated: ignition and fire growth (Cell2Fire) models. A cell i is selected for ignition using an ignition model such as random selection, custom probability function, or machine learning models that capture the risk of occurrence in the area. Once a cell is ignited, a weather scenario is selected and a wildfire is simulated. Each fire spreads through the cells following a provided fire-spread model and a shortest-path tree \mathcal{T}_i is obtained. After R replications, a multidigraph $\mathcal{GT}_{\mathcal{R}}$ is generated. Using this information, different risk maps are produced. (PHASE II) This module solves the problem of finding the adjacent cells that minimize expected losses due to future wildfires given specific constraints provided by the decision maker. Finally, the effectiveness of the treatment plans generated is evaluated.

Table 3.1: Basic description of the instances considered in this study including the total area in hectares, topographic characteristics, and dominant fuel behavior model type and codification following the Canadian FBP System for Dogrib, Arrowhead (AH), Revelstoke (RT), Mica Creek (MC), Glacier National park (GNP), Central Kootenay (CK), and Neptune Peak (NP).

Instance	Area [ha]	Elev. range [m]	Mean elev. [m]	Mean slope [%]	Dominant Fuel (Code)	Total fuels
Dogrib	79,611	[1299, 2825]	1693	21.04	Boreal Spruce (C-2)	11
AH	159,963	[430, 2894]	1329	42.96	Red/White Pine (C-5)	11
RT	350,956	[429, 2997]	1494	45.91	Red/White Pine (C-5)	11
MC	304,781	[516, 2895]	1523	48.41	Red/White Pine (C-5)	11
GNP	464,664	[300, 3000]	1701	45.79	Boreal Spruce (C-2)	11
CK	399,401	[300, 3000]	1811	51.64	Boreal Spruce (C-2)	11
NP	350,956	[429, 2997]	1494	45.91	Red/White Pine (C-5)	12

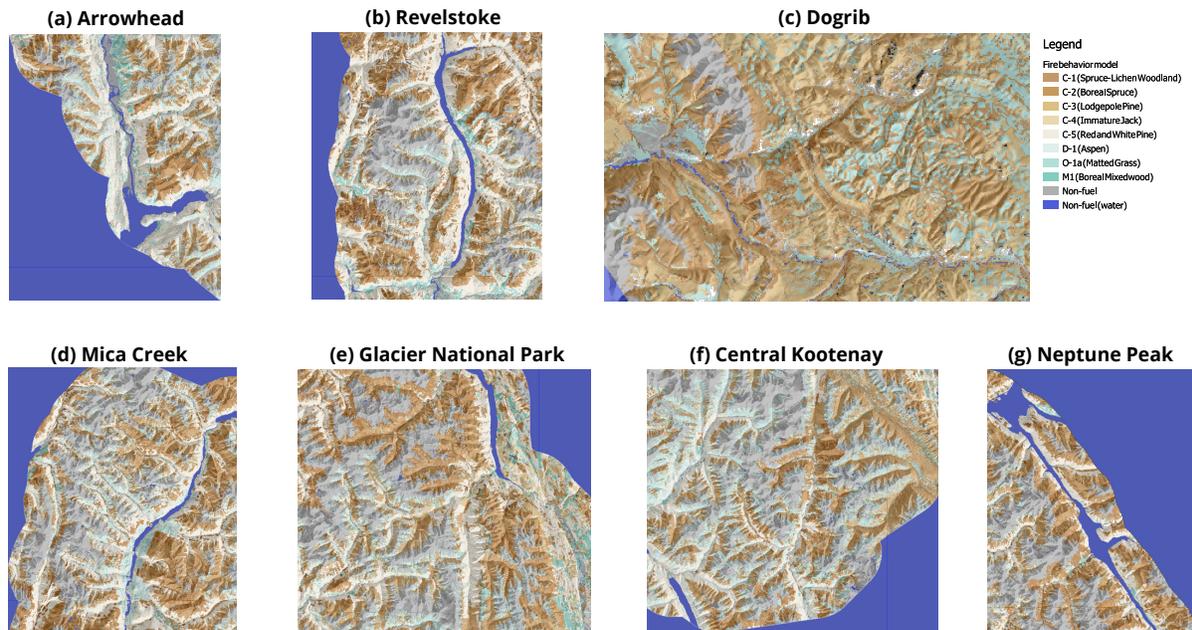


Figure 3.2: Map of the seven study areas. The legend indicates the color of the different fuel behavior models (fuel types) characterizing the terrain of each region. A hillshade effect has been applied to depict the elevation of the terrains.

Uncertainty

The location of an ignition point can be configured deterministically (DI) or randomly (RI). In a stochastic setting, ignition locations can be drawn from a given probability distribution, otherwise, a uniform distribution is set by default. The number of ignitions is set to be one per replication.

Weather scenarios include hourly observations of temperature, relative humidity, wind speed, wind direction, as well as their associated Fire Weather Index (FWI) System [345], fuel moisture codes, and fire behavior indexes from the FBP System [158]. The selection of weather scenarios can be deterministic (DW) or random (RW). In the first case, a predefined scenario is drawn. Otherwise, simulations are performed using a set of weather scenarios Φ , selected to represent multiple contexts in which a fire can spread, capturing the variability of this phenomenon. However, it is necessary to include only scenarios that boost the propagation of fire. As mentioned in [113, 120, 257] large fires are more useful to risk assessment. Historical data are obtained from the Climate Information Section of the Agriculture and Forestry site of Alberta, Canada, and data from the Yaha Tinda Auto station. This is the closest station to the coordinates of the forests used for simulation. We extract data of the daily averages of the weather parameters of the last 5 calendar years (2014 to 2018). From this data, we chose the 40 days with the worst weather conditions,

which favor the propagation of fire and proceeded to obtain $|\Phi| = 100$ scenarios of 24 hours.

Fire modeling

In this study we use Cell2Fire (Chapter 1), a fire growth model, presented in [250]. In its first version, this fire simulator has a similar performance to Prometheus (the most prominent simulator in Canada). Both consider the same aspects of a landscape such as: slope, aspect, and elevation; fuel types from the FBP System [125]; and weather to simulate fire growth. Cell2Fire is a cellular automata – fire spreads from cell to cell – unlike Prometheus which is based on wave propagation equations [339]. This is an advantage of Cell2Fire, because it can save and change the status of the cells in simulation time. During a simulation, a cell can be in one of five states: “*Available*”, “*Burning*”, “*Burned*”, “*Harvested*”, or “*Non-Fuel*” where the label “*Available*” indicates that the cell contains a flammable fuel type; “*Burning*” represents cells containing an active fire; “*Burned*” indicates that the fire has consumed the fuel available and passed through the cell; and “*Non-Fuel*” is a non-flammable fuel type such as rivers, lakes, or rocks. The “*Harvested*” state is provided so that the simulator can be embedded in a strategic harvesting planning system (see section 3.2). The harvesting planning module will be responsible for labeling the cells that are harvested and provide the appropriate post-harvest fuel type.

Detecting critical areas

We experiment and compare our new metric with three established approaches: burn probability maps, betweenness centrality, and fire protection value. The concept of burn probability (BP) maps is commonly used to indicate wildfire likelihood and assess the risk of fire by incorporating expected consequences [5, 7, 120, 239, 257]. Multiple software such as Burn-P3, FSim, and FlamMap [120, 257, 260] that calculate these metrics using different methods are available. Fuel treatment applications using this approach can be reviewed in [207, 258, 297]. Another metric applied in the context of fuel management is the betweenness centrality (BC) value [64]. It aims to identify and rank those areas that act as a link between sections of the forest. For example, in [145], the authors use this metric for the strategic placement of fuel treatments to mitigate the spread of fire. Other applications of this method in this context can be seen in [291]. In [253], the authors develop the fire protection value (FPV). Using this metric, the decision maker can identify sections of the forest where the fire tends to travel faster, focusing the treatment on these critical locations. The main difference with our approach is that the model is static in the sense that it relies on unique weather conditions, failing to capture temporal interactions. In order to capture the risk, these metrics (e.g., BP) are multiplied by expected consequences due to wildfire.

Downstream Protection Value

We model a forest as a two-dimensional lattice with identical cell areas. The attributes and spatial relation between these cells are modeled using an undirected planar graph $G = (\mathcal{N}, \mathcal{E})$ where the set \mathcal{N} represents the forest cells (graph nodes) and \mathcal{E} the edges between neighboring cells (see Figures 3.3 and 3.4). Nodes $i \in \mathcal{N}$ are associated with relevant attributes \vec{v}_i such as total volume available, selling price per cubic meter, treatment costs per area, fuel type, and any other relevant data provided by the user. Specific attributes can be considered for edges such as distances between cell centers, slope, transportation costs, among others.

When a fire occurs, a messaging process is generated between the nodes of G during the simulation, obtaining a directed graph $G_D = (\mathcal{N}_D, \mathcal{E}_D)$ where $\mathcal{N}_D \subseteq \mathcal{N}$ is the set containing all the cells burned during the replication and \mathcal{E}_D is built from these signals, representing the fire propagation between adjacent cells. Based on the network structure and the characteristics of the fire-growth model, we introduce the following notations/concepts:

- $Adj(i)$: denotes the set of available cells adjacent to $i \in \mathcal{N}$.
- $t(i)$: represents the time at which the fire reaches the center of cell i and starts propagating to its neighbors.
- We denote $i \rightarrow_t j$ if cell i is “**sending a message**” to cell j at time-step t , simulating and updating the fire evolution along the minimum distance between the center of both cells. Here, we assume that $t \leq t(j)$, the time to reach the center of cell j . When $t = t(j)$, the cell j changes its state to “*burning*” and we establish that “ **i sent a message to j** ” ($i \rightarrow_{t(j)} j$). Following this notation, we can represent the fire traveling time from the center of cell i to the center of cell j . Namely, $t_{(i,j)} = t(j) - t(i)$, and thus if $i \rightarrow_t j$ then $t(i) < t \leq t(j)$.
- For notation simplicity, we denote by $i \rightarrow j$ when the fire propagating from cell i reaches the center of cell j , i.e. $i \rightarrow_{t(j)} j$ occurs. Only the **shortest paths** are registered during the execution of the simulator.
- A cell i becomes inactive when no more messages can be sent from i . This occurs when no adjacent cells are available, $Adj(i) = \{\emptyset\}$. (e.g., surrounded by non-flammable or already burned cells).

We emphasize that the notation “ \rightarrow ” is only defined for adjacent cells on graph G , so in what follows, we extend its definition.

Definition 3 We say that “ **i reaches j** ”, and we denote it by $i \rightarrow j$ if there are $i_1, i_2, \dots, i_p \in \mathcal{N}$ such that:

$$i \rightarrow_{t(i_1)} i_1 \rightarrow_{t(i_2)} i_2 \rightarrow \dots \rightarrow i_p \rightarrow_{t(j)} j.$$

Note that the fire traveling time from i to j is given by:

$$(t(i_1) - t(i)) + \dots + (t(j) - t(i_k)) = t(j) - t(i).$$

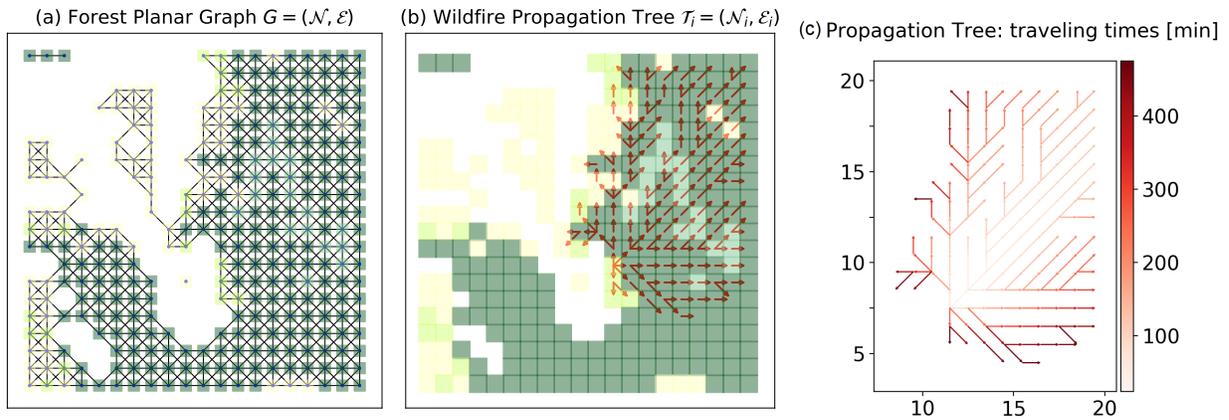


Figure 3.3: (a) Planar graph with edges indicating the connection between adjacent cells. (b) Simulated “shortest-path tree” from a heterogeneous instance with arrows indicating the propagation of fire during a replication. (c) Fire shortest traveling times. Red arrows indicate edges where fire was propagated (darker represents longer traveling time in minutes). Axes represent cell ids.

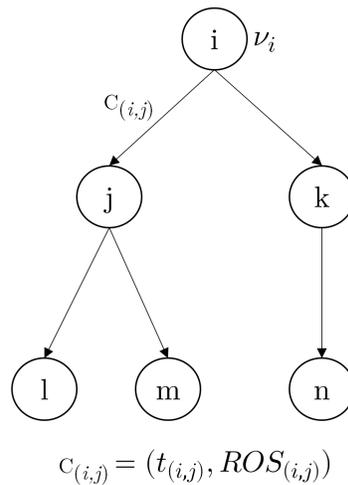


Figure 3.4: Example of a “shortest-path tree” graph generated from Cell2Fire. Cells are represented by nodes and the edges indicate that fire was propagated from one cell to another during the simulation. Weights $C_{(i,j)}$ include information w.r.t. the traveling times and average rate of spread (ROS).

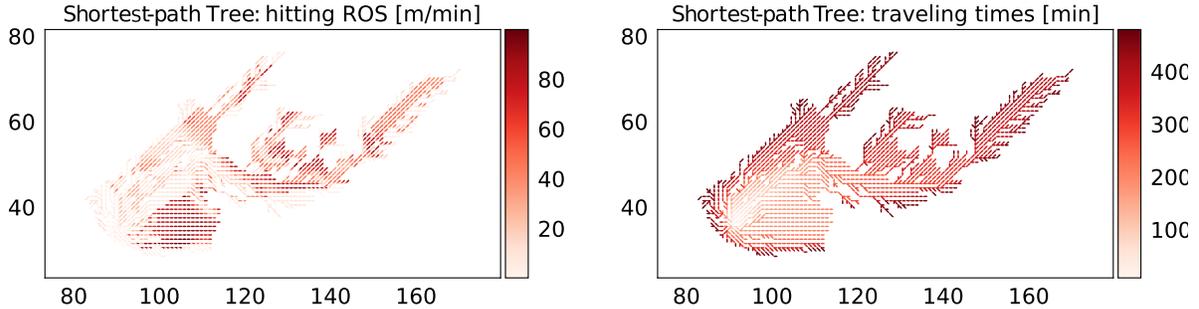


Figure 3.5: Shortest-path tree example from Dogrib instance. Arrows are used to indicate the propagation of fire between adjacent cells. Axes represent cell ids. Useful information regarding the ROS [m/min] (left) and travel times [min] (right) are registered at the spreading times between cells. Darker colors indicate larger values – faster rate of spread and longer traveling times, respectively.

Definition 4 We say that the graph $\mathcal{T}_i = (\mathcal{N}_i, \mathcal{E}_i)$ is the “shortest-path tree” of the cell i if:

- $i \in \mathcal{N}_i$ is the root of the tree \mathcal{T}_i , where $\mathcal{N}_i = \{j : i \rightarrow j\} \cup \{i\}$.
- $\mathcal{E}_i = \{(k, j) : k, j \in \mathcal{N}_i, k \leftrightarrow j\}$ where \leftrightarrow represents the existence of an edge between cells k, j due to the propagation of a fire message during the simulated wildfire.
- Each edge $e = (k, j) \in \mathcal{E}_i$ is associated with a weight $C_{(k,j)} = (t_{(k,j)}, ROS_{(k,j)})$ including the traveling time between cells and average ROS during the new ignition.

The concept of “shortest-path tree”, although simple, is very powerful and useful. Given any cell inside the forest, it summarizes the fire propagation dynamics during a simulation, starting from this particular root cell. Intuitively, we would like to decrease the size of each “shortest-path tree”, meaning that the number of cells burned after a wildfire is minimized. We can easily obtain useful insights and assess the risk of a cell during a fire season by noting the size of its tree. The bigger \mathcal{T} , the larger the number of cells burned because of the presence of the root cell, acting as the shortest link between two different areas.

Traveling times recorded inside the “shortest-path trees” provide useful information regarding the shortest propagation paths inside the forest. From the discussion above, each edge (i, j) represents the shortest path of the fire from the ignition point to cell i and from i to j . Furthermore, average propagation time and ROS can be calculated for a tree \mathcal{T}_i , providing practical insights when developing the fuel treatment plan by also incorporating relevant elements for operational fire suppression (e.g., average propagation speed) that can be critical in the event of a future wildfire. From these outputs, interesting conclusions can be extracted with respect to the role of different fuel types inside the propagation dynamic,

as well as weather conditions that affect the effective ROS between cells, among others (see Figure 3.5).

Although focusing the attention on the shortest propagation paths of the fire within the landscapes like the minimum travel time (MTT) algorithm developed in [118], our method presents significant differences both at conceptual and methodological levels. First, the MTT algorithm is based on the calculation of the shortest paths following the wavefront propagation equations developed by [280] while our fire growth algorithm is based on a cellular automata model following the Canadian FBP system [156] equations. Second, while fire perimeters are calculated following an interpolation approach with the MTT algorithm, these are directly generated from the fire spread and growth dynamic in our method as we register the fire progress at a given time frequency, thus obtaining fire perimeters directly from the equations governing the fire spread model. Third, although both algorithms can be used in a stochastic setting (e.g., uncertain weather) the original MTT fire growth algorithm calculates and estimates fire growth in the absence of time-varying conditions (wind direction, wind speed, moisture content, etc.) focusing on the fuel and topographic characteristics of the landscape, given initial weather conditions. In contrast, our approach is designed to integrate both constant and dynamic conditions. Following this line, the nature of the original MTT algorithm is deterministic, leading to the exact same outputs given a set of inputs. However, alternative procedures can be applied to introduce uncertainties in its execution. Our shortest-path trees can incorporate stochasticity via (1) probabilistic weather conditions following a given distribution for all different components (at any temporal granularity) and (2) allowing the perturbation of the ROS calculations via a coefficient of variation to account for approximation errors of the fire spread model, potentially leading to multiple outputs given the same initial conditions.

We introduce the concept of “Global Propagation Tree”, the graph generated from the superposition of individual trees (Figure 3.6). In the context of multiple replications – wildfire simulations under different conditions – this graph allows us to calculate all the required parameters for our optimization model such as ignition probabilities, expected area burned, average rate of spread, and risk metrics.

Definition 5 We define the graph $\mathcal{GT}_R = (\mathcal{N}_R, \mathcal{E}_R)$ as the “*Global Propagation Tree*” of the landscape after R replications where we have:

- The set $\mathcal{N}_R := \cup_{r \in R} \mathcal{N}_r$ contains all the ignited/burned nodes among all replications R .
- $\mathcal{E}_R := \cup_{r \in R} \mathcal{E}_r$, the set of all edges traversed by fire during the R replications.
- The weights $C_{(i,j),R} := (t_{(i,j),R}, ROS_{(i,j),R}, f_{(i,j),R})$ associated with any edge $(i, j) \in \mathcal{E}_R$ where:

$$- t_{(i,j),R} = \sum_{r \in R} \frac{1_{(i,j) \in \mathcal{E}_r} t_{(i,j),r}}{|R|} \text{ the average traveling time between node } i \text{ to } j.$$

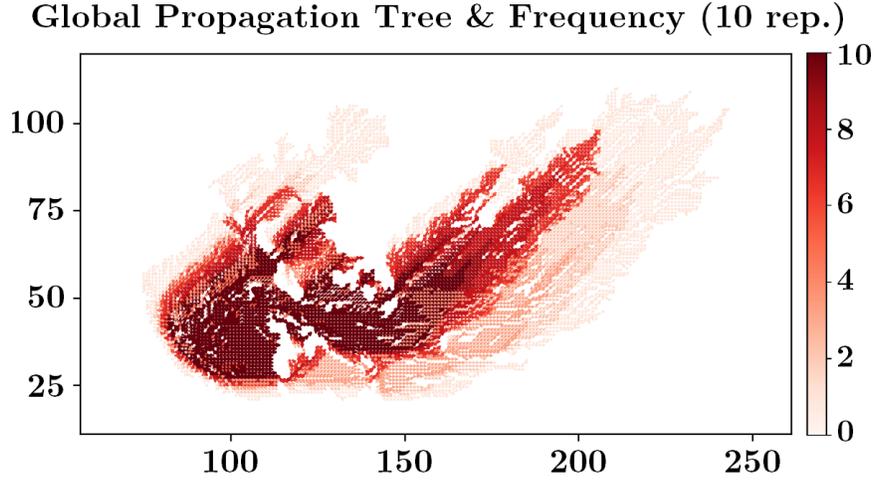


Figure 3.6: Global Propagation Tree obtained from the overlapping of 10 replications (individual shortest-path trees) starting from the same ignition point including uncertainty in the weather scenario $\phi \in \Phi$ in Cell2Fire. Edges frequency usage among replications is indicated with colors, where darker tones represent common propagation patterns. Axes represent cell ids.

- $ROS_{(i,j),R} = \sum_{r \in R} \frac{1_{(i,j) \in \mathcal{E}_r} ROS_{(i,j),r}}{|R|}$ the average ROS when fire hits node j from i .
- $f_{(i,j),R} = \sum_{r \in R} 1_{(i,j) \in \mathcal{E}_r}$, the number of times an edge (i,j) is traversed by fire among the $|R|$ replications.

The construction of Global Propagation Trees will be used to generate outputs and metrics when performing multiple replications accounting for a diverse set of potential wildfires, representing another difference between our method and existing algorithms such as the mentioned MTT where individual propagation paths are processed independently and not combined into a multidigraph to weight relevant metrics.

Mathematical Model

Following the idea of evaluating the level of protection given by treating a certain area, we introduce the downstream protection value (DPV), a new fire risk model that aims to measure the value of what is burned downstream, starting from a burning cell i . Thanks to this metric, we can quantify the impact of treating cell i in terms of an optimization problem defined by the decision maker, optimizing the fuel treatment plan when planning under wildfire uncertainty.

DPV formula	G_1	G_2	G_3
$\sum_{j \in \mathcal{N}_i} V_j$	6	6	6
$\sum_{j \in \mathcal{N}_i} d(i)V_j$	12	24	6

Table 3.2: DPV comparison example.

Definition 6 Let $i \in \mathcal{N}$, and $\mathcal{T}_i = (\mathcal{N}_i, \mathcal{E}_i)$ its shortest-path tree with root node i . Then, the Downstream Protection Value, $DPV(i)$ is defined as

$$DPV(i) = \sum_{j \in \mathcal{N}_i} NV_j$$

where NV_j is an appropriate value for node j .

From the previous definition, DPV is the summation of “node values” (NV) factors, defined as a common value unit used to represent the relative importance of each cell and prioritize the treatment within the landscape as we seek to minimize the impact of (e.g., expected losses) wildfires. For example, the total volume per hectare arises as a consistent and simple NV variable. We can easily extend it and include more information by weighting values by probabilities and other characteristics of the nodes/edges (e.g., the degree of the node). In addition, due to the additive property of the model, we notice its natural extensions to work with larger units (cluster of cells) and to calculate the expected impact per cell from multiple replications R , using the Global Propagation Tree \mathcal{GT}_R instead of \mathcal{T} .

In order to illustrate potential applications and limitations depending on the selected NV, we proceed to analyze and test different possibilities.

Suppose that we have three different “shortest-path trees” (Figure 3.7) and we want to calculate the DPV using just the volume as our NV. For simplicity, we assume identical cells (1 unit of volume). In this case, the DPV is exactly the same for cell 1 in all graphs, but their structures, traveling times, and ROS are significantly different, therefore, this NV may not be appropriate for the situation.

We tested different extensions and modifications to our initial DPV: (1) Adding a decay factor $\alpha(j)$ depending on the depth inside the tree, (2) weighting the NV by average traveling times $t(j)$ or ROS, and (3) multiplying the original summation by the degree $d(i)$ of the node. In Table 3.2 and Figure 3.8, we can see how different values can be obtained depending on the approach used. Weighting the volume by the total degree $d(i)$ of the node allows the DPV to better capture this information, increasing the metric value when cells tend to be more connected inside the network, and thus, tend to spread fire to more sections of the generated tree.

This way, our model extends previous studies and metrics such as the node influence grid concept [114], where a grid summarizing the length of the propagation paths for each cell

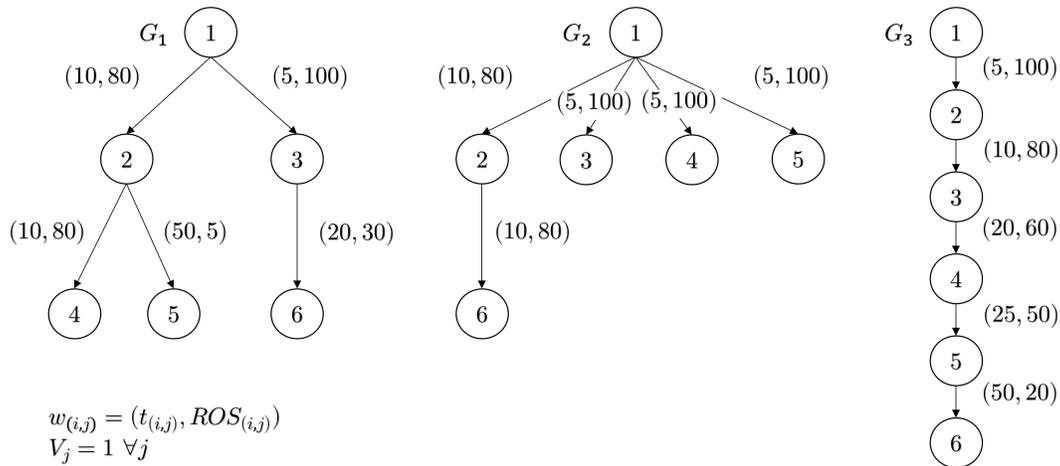


Figure 3.7: DPV comparison for three different “shortest-path tree” graph structures.

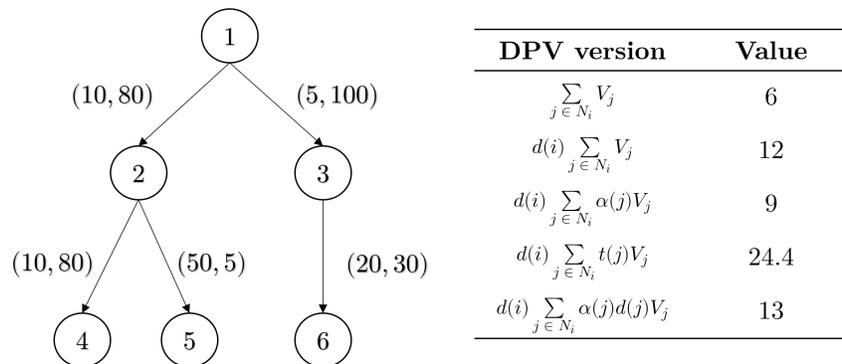


Figure 3.8: DPV versions example. Several combinations of relevant variables of the graph can be used for calculating the node values within the landscape. Its components will depend on the objective and planning horizon of the study.

– not accounting for differences between nodes characteristics – is generated. Here, instead of calculating metrics by averaging values from individual shortest propagation paths after multiple replications, we condense all shortest-path trees into a unique graph, Global Propagation Tree. As discussed, this new graph is a multidigraph summarizing all individual shortest propagation paths, indicating the frequency each edge was used across all replications. This allows weighting the relevant output metrics (e.g., DPV) by the degree and characteristics of each node. Although a subtle difference, this approach leads to significant differences in the outputs even in small instances, as shown in Figure 3.8 and Table 3.2.

Optimal fuel-treatment plans

We formulate the fuel-treatment problem as an optimization model with the objective to find the subset of adjacent cells that maximize the total summation of the NVs considered, e.g., that maximize the disruption of the shortest propagation paths. Depending on the objective of the treatment plan, this value function could incorporate the impact of a unique fire behavior output (e.g., burnt probability, ROS and traveling times, fire intensity, flame length, etc.) or general metrics such as DPV or BC including the relevant assets of the landscape (e.g., timber volume). In this work, we use a version of the DPV metric where the volume available in each cell i of the landscape is weighted by its corresponding degree $d(i)$ (Definition 4). This parameter is extracted from the Global Propagation Tree obtained after $|R|$ independent replications.

The formulation allows multiple constraints including resource allocation (e.g., total budget), demand for forestry products, ecologic/demographic restrictions (e.g., protected areas), among others. In this work, we fix the total number of nodes that can be treated, varying the total treatment fraction to compare plans with different intensity instead of enforcing an explicit budget constraint. Adjacency constraints are included with the purpose of obtaining more realistic and easily adaptable solutions in practice [75], where the fuel treatment is carried out in connected patches. Formally, we solve the following general optimization problem:

Definition 7 *Given a weight $w_i \geq 0 \forall i \in \mathcal{N}$ and an undirected graph $G = (\mathcal{N}, \mathcal{E})$, we aim to maximize $\sum_{i \in S} w_i$ by selecting the subset of nodes Δ inducing a connected sub-graph G_Δ . A set of operational constraints (e.g., treatment resources) define a feasible region \mathcal{F} . We formulate the Protection Value Problem (PVP) as:*

$$\begin{aligned}
 (PVP) \quad & \max_{\Delta \in \mathcal{G}_\mathcal{N}} \sum_{i \in \Delta} w_i \\
 & \text{s.t. } \Delta \text{ connected} \\
 & \Delta \in \mathcal{F}
 \end{aligned} \tag{3.1}$$

where $\mathcal{G}_\mathcal{N}$ is the set of all subsets of nodes in \mathcal{N} , and w_i is the value of each cell $i \in \mathcal{N}$ (e.g., obtained from BP, BC, FPV, or DPV).

Therefore, for a given undirected and node-weighted graph $G = (\mathcal{N}, \mathcal{E})$, we want to find the subset of nodes with maximal sum of DPV (or another relevant metric), inducing a connected subgraph Δ . This is a variant of the maximum-weight connected subgraph problem (MWCSP) a known NP-hard combinatorial problem [21, 166], where extra constraints are included.

In this work, we solve this optimization problem in two steps. (1) A warm start is provided by a greedy heuristic that selects which cells to harvest as follows. After ranking the cells by their weights w_i , it adds adjacent cells to the current patch with the highest w_i value until a stop condition (e.g., maximum number of treated cells) is satisfied. (2) Then, we solve a Mixed Integer Programming (MIP) based on a network flow structure as in [97]. This exact model uses the feasible solution provided by the heuristic as a starting point for the branch and bound method, improving its convergence. Let x_i be a binary variable for each cell $i \in \mathcal{N}$, indicating if cell i is included (1, selected for treatment) in the connected subgraph \mathcal{G}_Δ or not (0). Each cell is characterized by an NV, w_i , obtained after R replications, a certain cost c_i , and a utility u_i for performing the fuel-treatment activities (e.g., harvesting cost and timber selling price). We define a continuous variable $y_{i,j}$ for each edge $(i, j) \in \tilde{\mathcal{E}}$, the set of directed edges obtained from the original set \mathcal{E} . These variables represent the total flow from cell i to cell j in the network formulation, allowing us to model the connectivity constraints by forcing a continuous flow among the selected units.

To model the number of nodes that can be selected, an extra cell s acting as the source of the flow is introduced to the network obtaining $\mathcal{N}_s = \mathcal{N} \cup \{s\}$, the set of all nodes in the network. We connect this node to each cell inside the original set \mathcal{N} using the directed edges of the set $\mathcal{E}_s = \tilde{\mathcal{E}} \cup \{(s, i) \mid i \in \mathcal{N}\}$. All the flow from the source s is then fed into the original network and/or absorbed by an auxiliary variable $z \in [0, d]$ with d the maximum number of cells to be treated (eq. 3.3). This, to allow solutions selecting less than d nodes. Adjacency constraints, on the other hand, are enforced by eqs. (3.4)-(3.7). In eq. (3.4), at most one cell $i \in \mathcal{N}$ is acting as a link between the source s and the original network using the binary variable α_i , equal to 1 if the cell i gets the flow from the source. The remaining cells will not get any flow from the source, setting those flow variables to zero (3.5). Constraint (3.6) ensures that cells included as part of the solution retain a positive flow value (one unit), forcing their selection. Then, we model the conservation of the flow inside the network (3.7). Here, each cell $i \in \mathcal{N}$ with a positive flow acts as a sink, consuming one unit of flow (i.e.,

part of the solution) and any remaining flow is sent to adjacent cells l if an edge (j, l) exists.

$$(PVP) \quad \max \sum_{i \in \mathcal{N}} w_i x_i \quad (3.2)$$

$$s.t. \quad z + \sum_{j \in \mathcal{N}} y_{s,j} = d \quad (3.3)$$

$$\sum_{i \in \mathcal{N}} \alpha_i \leq 1 \quad (3.4)$$

$$y_{s,j} \leq d \alpha_j \quad \forall j \in \mathcal{N} \quad (3.5)$$

$$y_{i,j} \leq d x_j \quad \forall (i, j) \in \mathcal{E}_s \quad (3.6)$$

$$\sum_{i \in \mathcal{N}: (i,j) \in \mathcal{E}_s} y_{i,j} = \sum_{l \in \mathcal{N}: (j,l) \in \tilde{\mathcal{E}}} y_{j,l} + x_j \quad (3.7)$$

$$\forall j \in \mathcal{N}_s$$

$$\sum_{i \in \mathcal{N}} u_i x_i = U \quad (3.8)$$

$$\sum_{i \in \mathcal{N}} c_i x_i = C \quad (3.9)$$

$$x_i, \alpha_i \in \{0, 1\} \quad \forall i \in \mathcal{N} \quad (3.10)$$

$$y_{i,j} \in \mathbb{R}^+ \quad \forall (i, j) \in \mathcal{E}_s \quad (3.11)$$

$$z \in \mathbb{R}^+ \quad (3.12)$$

Equations (3.8) and (3.9) are included in the (PVP) model in order to keep track of the total utility (U) and cost (C) of the optimal treatment plan (if provided). Notice that we can also compute them after solving the model, however, we include them in the formulation as an example of further extensions of the original model, where extra constraints (e.g., maximum budget) can be included.

If needed, we can force the fuel-treatment plan to include the desired amount of cells d to be treated adding the following equation:

$$\sum_{i \in \mathcal{N}} x_i = d \quad (3.13)$$

This way, the optimal connected subgraph S will include exactly d cells. Notice that the same effect can be achieved by setting $z = 0$. In addition, we can include eq. (3.14) as an extra cut to ensure that the total number of selected cells is equal to the total flow inside the network.

$$\sum_{i \in \mathcal{N}} x_i = \sum_{j \in \mathcal{N}} y_{s,j} \quad (3.14)$$

This formulation of the (PVP) can be solved directly using state-of-the-art solvers such as Gurobi or CPLEX when the size of the instance is relatively small. However, obtaining

a high-quality solution (below 2% of gap w.r.t. the best bound) or even a feasible solution within reasonable running times becomes a challenge when dealing with large landscapes or high-resolution instances with a significant number of cells. In order to address this situation, we implement the two-step method described above where the solution obtained from a greedy heuristic, selecting adjacent cells in decreasing value of w_i , is used as a warm-start for the exact model. This allows us to obtain low final gap solutions – an average below 2% – in reasonable solving times (less than 5 minutes when solving the largest and densest instances using a daily laptop, less than 1 minute for the smaller instances) across all instances.

3.3 Results and Discussion

A total of 65,500 experiments, involving testing the different metrics with (i) heterogeneous landscapes of varying sizes – 79,611 to 559,746 cells – and land cover compositions, (ii) different numbers of replications $R \in \{10, 100, 1000\}$, (iii) a random or deterministic selection of the weather scenarios and ignition point per replication, and (iv) varying the total fraction tf representing the maximum area of the landscape to be treated given the decision maker’s resources are conducted. Due to size limitations, fires are simulated for 12-hour windows in the Alberta province, while full-day scenarios are used in the British Columbia instances (see Appendix C for detailed results). For any given wildfire, using our model decreases the average rate of spread and expected area burned by 53% and 57%, respectively, compared to its closest competitor (BC) thus creating effective fire-resilient landscapes. These improvements are observed across all instances and weather scenarios, reaching differences above 60% of the total area protected with respect to the second-best metric when treating 5%, 10%, and 15% of the landscape. The results are useful when facing budget and operational limitations, where resources must be spent in the most economical manner. The DPV performs best in all experiments, rapidly extracting the core fire propagation patterns within the landscape. It identifies high-risk regions, focusing the treatment on those zones, with its effectiveness elevated as the number of replications is increased. We measure this amelioration by noting the significant decrease in expected area burned – with differences up to 67% of the total area protected – when compared with the second-best metric (BC) across all treatment levels.

In general, we observe that increasing the number of replications positively influences the performance of all methods by decreasing (i) the volatility of the results and (ii) the total expected area burned after the application of the fuel-treatment plan. These effects are evident in the RW-RI scenarios, where uncertainty is a crucial part of simulated fires. DPV not only surpasses the front line alternatives but it has the fastest learning capacity with an improvement up to 65% in the protected area when increasing R from 10 to 100.

However, this positive relation is not true for all metrics. This can be seen in the Arrowhead instance (Figure C.2, top-center), where including 100 replications does not improve BP treatment plan’s effectiveness, yielding a global average for $tf \leq 25\%$ of 1066.7 ± 1191.7 burned cells when training the metric with 100 replications (BP 100) versus 1099.6 ± 1090.3

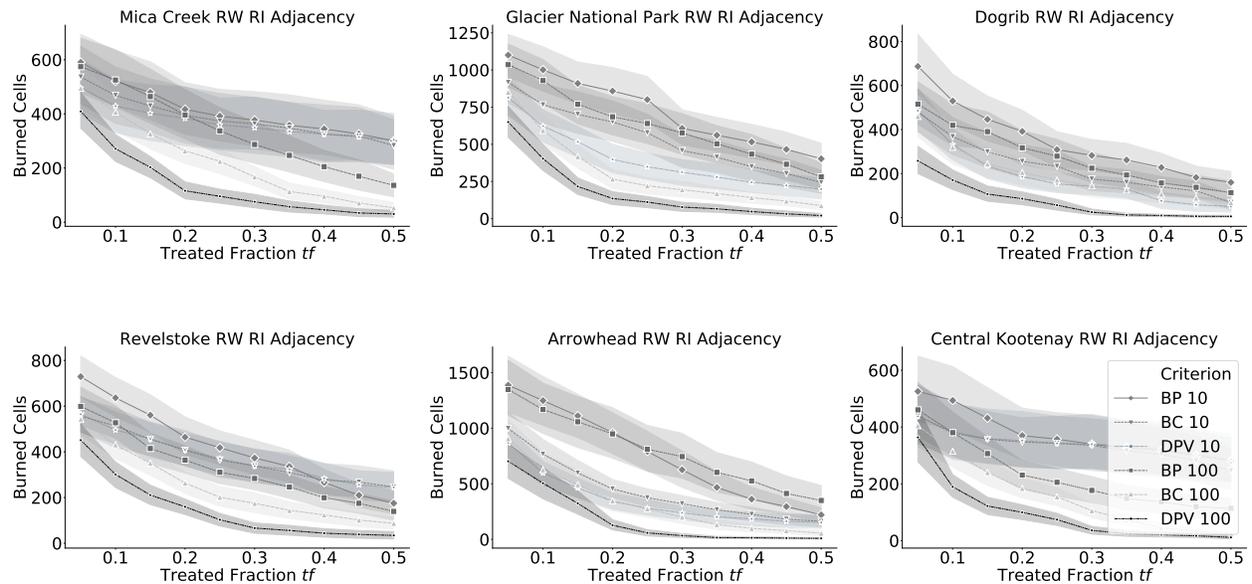


Figure 3.9: Multiple treatment plans are tested for different fraction levels (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications e.g., DPV 10, represent the standard deviation of the treatment performance after evaluating the new landscape with 100 independent replications. Adding new simulations from potential wildfires significantly improves the performance of our metric, which is critical to generate effective treatment plans. As an example, focusing on the Arrowhead instance $tf = 0.25$, we observe that using the DPV model, we decrease the average area burned from 281.1 ± 348.0 (DPV 10) to 58.4 ± 139.1 ha. (DPV 100). Here, we notice that even training the DPV model with 10 samples is almost as good as those of BC 100 (second-best), showing the potential of our method. Increasing R , DPV is able to significantly improve its performance, learning faster than any other metric reaching up to 75% of improvement when more data are provided.

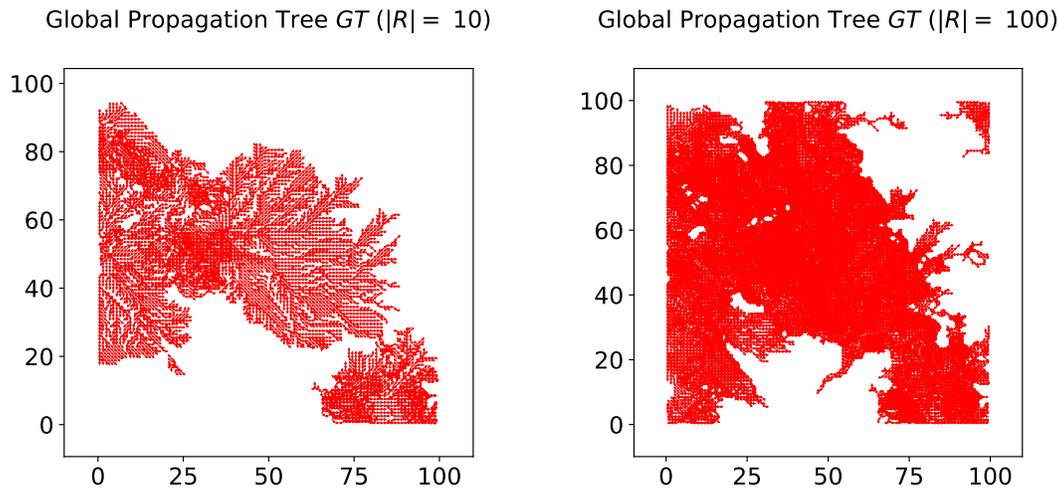


Figure 3.10: “Global Propagation Trees” for $|R| \in \{10, 100\}$ replications of the Dogrib RW-RI instance obtained after the overlapping of all individual shortest-path trees in a 100×100 patch (axes). More complex and general wildfire patterns are captured when increasing the number of replications (right) as larger – unexplored with fewer replications – areas of the landscape are covered, allowing the metrics to exploit this information when ranking the cells for the treatment plan.

when including only 10 training samples (BP 10). We omit the results of the FPV model since it is consistently outperformed by the DPV/BC both in the protection of the landscape and in the computational performance.

We can illustrate the previous discussion by observing the “Global Propagation Trees” \mathcal{GT}_R obtained from a section of the Dogrib instance for a different number of replications R (Figure 3.10). From the plots, the number of potential trajectories captured when including more simulations is clear: \mathcal{GT}_{10} contains limited information for all nodes in the graph, leading to classify several cells with null DPV (or alternative metric) since no observation include them into the wildfire propagation dynamic. On the other hand, the \mathcal{GT}_{100} graph provides a representative and exhaustive picture of the potential fire propagation paths in different sections of the landscape, allowing to easily identify which sections should be treated, as well as gathering relevant information regarding the expected ROS and traveling times of the fire within the forest.

Comparing the impact of the different sources of uncertainty included in the experiments (Figure 3.11), we observe similar trends: the DPV is able to converge faster toward full protection of the landscape when fixing the ignition area (DI) generating the most prepared landscapes for addressing fires under weather and ignition (RW-RI) uncertainty across all tested instances.

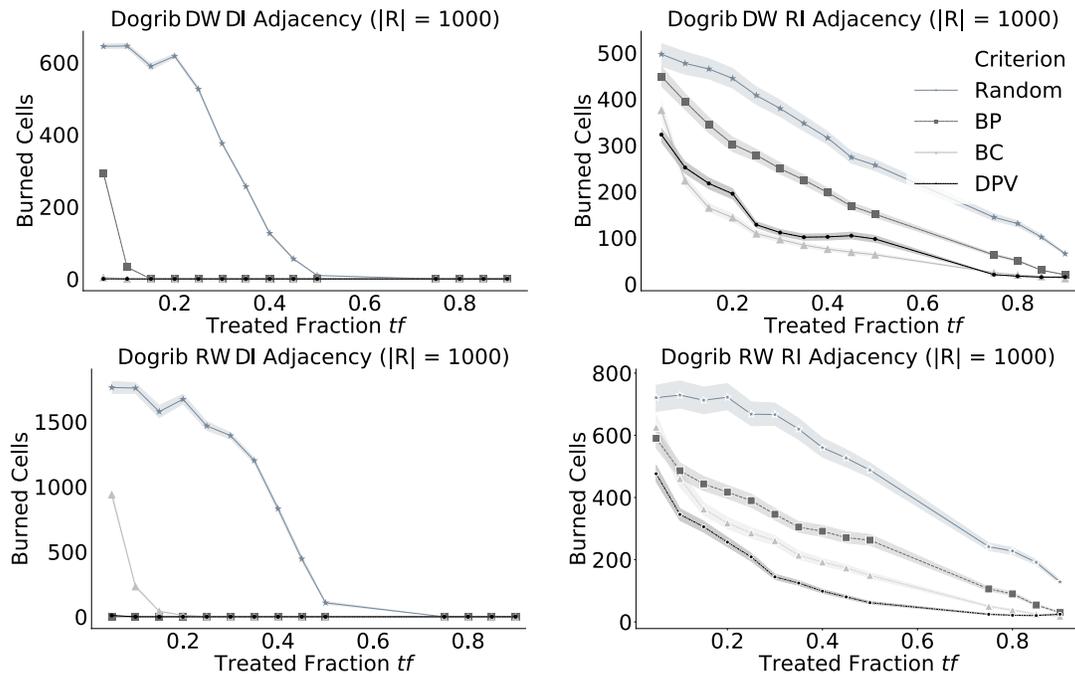


Figure 3.11: Treatment plan effectiveness (y-axis) depending on the simulated scenario (ignition point and weather, e.g., DI-DW) and treatment level (x-axis) for Dogrib instance. Fixing the ignition point (DI) to a bounded area (left side) significantly simplifies the problem, obtaining fast convergence toward no losses due to wildfire with all metrics.

Our DSS could be implemented in several practical applications. As an example, it could capture the fire propagation patterns for a particular forest by replicating historical wildfires and their weather conditions, obtaining a fully adapted and trained model. Moreover, its data-driven approach allows decision makers to continuously update the parameters of the optimization model, allowing them to add new information from the landscape and incorporate it into the optimal fuel-treatment plan. Different subsets of the forest could have different values (e.g., see Figure E.5). Factors such as native forest areas, highly population-dense sectors, or regions naturally predisposed for hosting certain species should be accounted for in practice. These can be easily incorporated in our framework by modifying the NV included in the DPV model as well as adding specific constraints to the *PVP* model. Thus, certain sections of the forest could have different weights when optimizing the plan. This difference is translated into multiple DPV heatmaps, and thus, different treatment plans are obtained, showing the flexibility of the framework.

A series of valuable outputs are obtained from the execution of the proposed methodology, allowing the modification of relevant values associated with the nodes based on the information gathered from the experiments. For example: (1) ROS heatmaps projected on

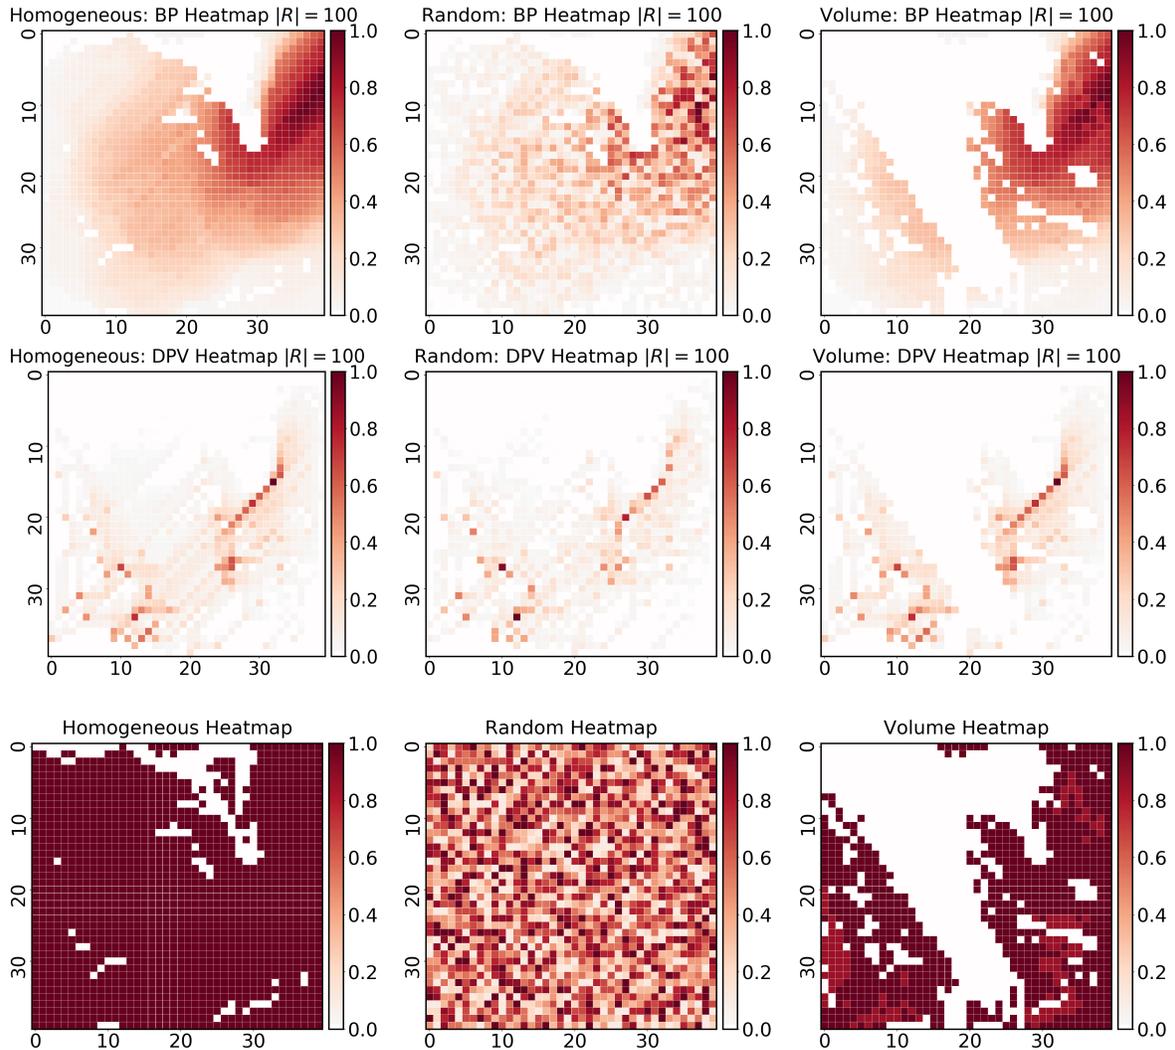


Figure 3.12: Different heatmaps example. Heatmaps obtained after a hundred replications using BP (first row) and DPV (second row) metrics for three different landscape – normalized – value functions (third row). A homogeneous (i.e., all flammable cells with identical value), random (each cell with a random uniform value between 0 and 1), and a volume-oriented value function are used to illustrate the impact on the heatmaps obtained by the metrics, leading to significantly different treatment plans.

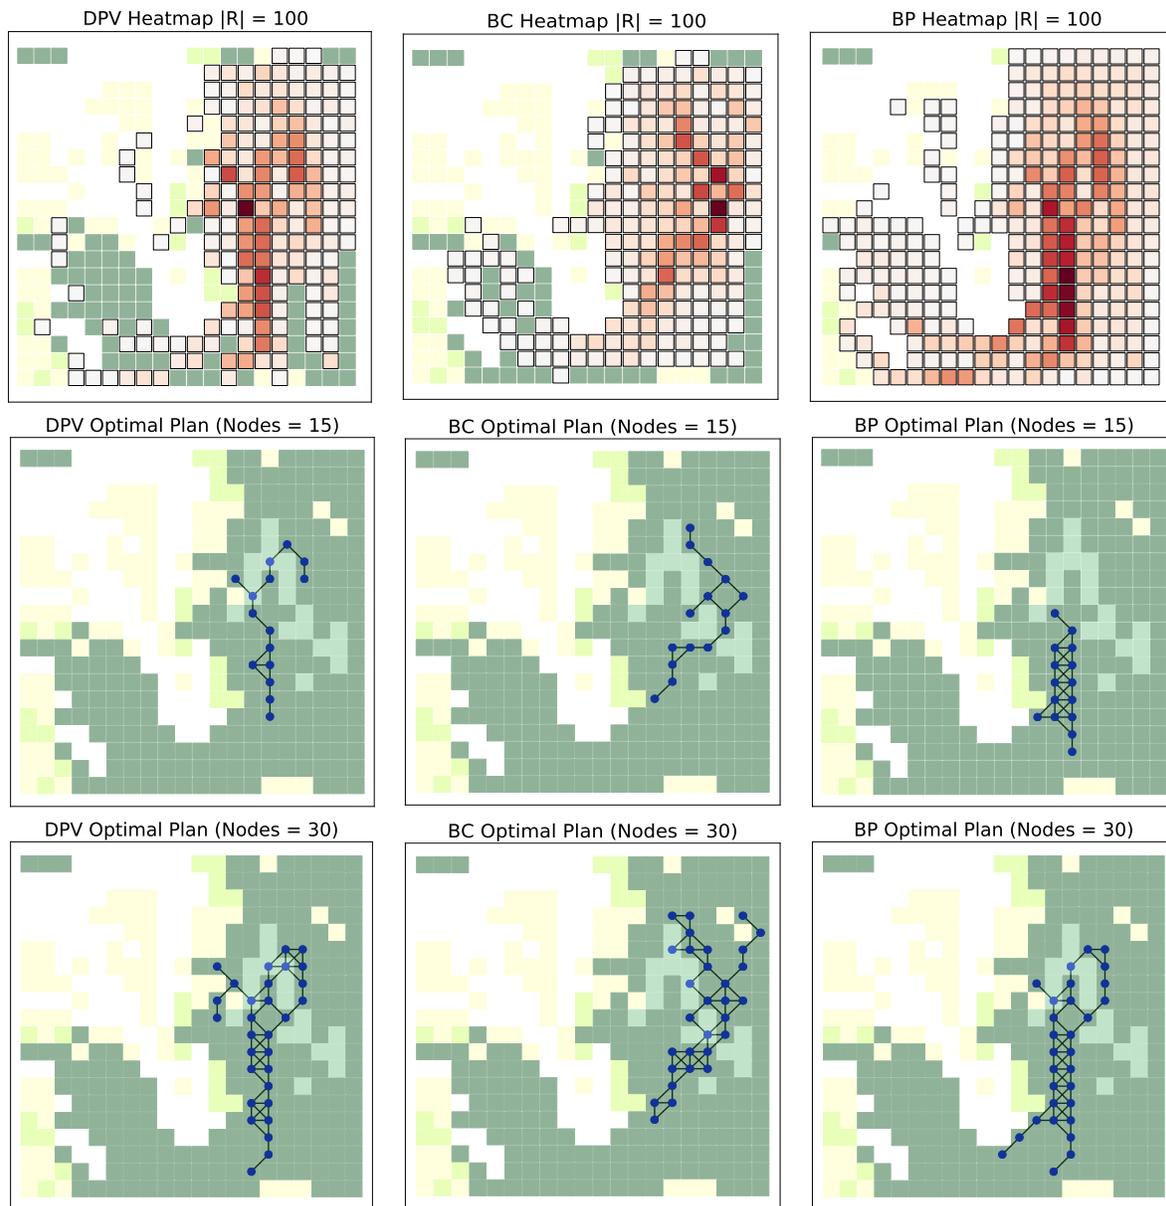


Figure 3.13: Full pipeline example. Heatmaps obtained after a hundred replications using DPV, BC, and BP metrics. Blank cells represent non-fuel sections of the landscape. Highlighted and darker red cells indicate higher metric values (first row). Values are then fed into the protection value problem (PVP), obtaining optimal treatment plans selecting the continuous patch of cells that minimizes the expected wildfire losses (dots connected by edges) for each different metric and treatment fraction level tf such as 5% (second row) and 10% (third row).

the “shortest-path trees” can be analyzed to determine which sections of the landscape tend to rapidly spread the fire under certain weather scenarios. Thus, helping the decision maker to identify highly flammable and risky land covers within the forest as well as the most frequent fire propagation lines. (2) Shortest traveling times between all nodes inside the “Global Propagation Tree” could provide critical information when planning the suppression activities and resource allocation. Knowing the frequency and propagation patterns of the most likely potential fires within the landscape could be used to generate efficient prevention plans and to estimate the number of units of different suppression resources (e.g. helicopters, drones, firefighters) needed to act under several scenarios, in order to suppress a detected ignition. (3) Custom metrics based on the DPV model can be easily tested and adapted by the user. After developing a series of experiments for a landscape, researchers could customize the metric by adjusting or even proposing a new value for each cell, exploiting the knowledge gathered from multiple replications and the performance achieved by the proposed metrics during the experiments. This way, more general and complex metrics – including elements such as lightning strikes probabilities – can be easily tested in our open-source DSS, giving the user full flexibility and expanding the potential of the system.

The proposed model can take advantage of historical data to develop efficient fuel-treatment plans with our DSS, obtaining landscapes that are better prepared for dealing with wildfires in comparison to state-of-the-art metrics. Overall, DPV emerges as an effective option for minimizing large wildfires with significant room for improvement thanks to its flexible and expandable model. Fuel-treatment plans generated from our model are able to detect high-risk areas, focusing the treatment on the most critical sections of the forest. Therefore, they result in optimized forest structures that minimize wildfire impact (see Figure 3.13).

Compared to previous studies, the degree of the node is crucial to calculate its importance (DPV). This is a fundamental difference with respect to similar metrics such as the node influence grid value [114] or BP as we take into account not only the role of the node regarding the depth but the width of the fire propagation within a landscape. Larger degree values indicate that one cell tends to disseminate the fire to a broader area of the landscape whereas smaller values indicate that it tends to propagate the fire in a unique direction. As mentioned in section 3.2, this avoids bias toward long/deep unidirectional chains of nodes as well as increasing the relative importance of those nodes with a deeper and wider impact on the fire propagation. This is translated into more effective plans, decreasing the average expected burned area in our experiments approximately by 20% (on average) compared to the non-weighted approach. Moreover, node influence grids do not take into account the characteristics of the cells, focusing only on the length of the shortest paths. In our method, the incidence of each node is weighted by a certain value (e.g., the volume of timber available) that represents the relevant characteristics of the cell – *it is not the same to burn grass or trees* – with respect to the optimization objective, leading to different treatment plans.

Similarly, DPV calculations are performed in a unique multidigraph summarizing all individual shortest propagation paths. This makes it efficient and fast to calculate as only a unique graph is processed instead of calculating individual metrics for each replication (thus

processing R graphs).

Although the observed results highlight the theoretical effectiveness of the proposed method, they should be interpreted with caution for practical applications. A series of simplifications and assumptions must be covered in future research to address more complex instances. In this work, we covered one (treatment intensity) of the main five factors affecting the fuel treatment design and optimization. Further research is needed to address the impact of treatment unit size, shape, and spatial distribution. The model can be extended in these directions, evaluating the performance of the metrics at different aggregation levels, allocations, and fuel treatment distribution. In the same line, we aim to extend the number of actions available in the model. Currently, only one simplified fuel treatment action (mechanical treatments transforming flammable cells into non-flammable ones) is included. This is clearly a limitation for practical applications where a variety of actions are available to treat a cell, with potential different objectives.

The current framework is static. Every time we select a cell to treat, we are modifying the protection values of other cells, thus perturbing the values of the network. This problem is significantly harder and more complex than the current version, requiring further modeling efforts. Similarly, it can be extended to a multistage setting assessing the effectiveness of iterative fuel treatment efforts within a given planning horizon. Moreover, the results presented in this work are conditional upon the occurrence of fires within specific time frames, meaning that the evaluation of the effectiveness of the fuel treatment plans assumes that new wildfires occur during the period of time in which treatments are highly effective. This may not be the case in practice, where the temporal dimension (when to apply the treatment) is crucial to implement effective plans. In addition, future iterations of the framework should focus on modeling realistic budget constraints, allowing us to obtain practical fuel treatment plans. We can represent the inherent trade-offs faced by decision makers when selecting the optimal plan by including multiple objectives, thereby providing relevant managerial insights. Following this line, the next Chapter will cover an extension of this framework including multiple layers that could be incorporated as part of the weights of the nodes when calculating their importance within the network. This, to provide a more general and robust framework, potentially leading to alternative solutions incorporating the impact of wildfire and the treatment effects on ecosystem services, wildlife habitat, populated areas, among other relevant layers of information.

3.4 Conclusions

There is strong evidence from regions across the world that the warming and drying of landscapes due to climate change are connected to increasing fire occurrence, intensity, and duration. Studies have shown that globally, the rise in fire occurrence, intensity, and duration are related to environmental effects caused by climate change. As a result, current fire suppression activities, resources, and fire management efforts are not enough to protect forests without an optimized structure to decrease the impact of future wildfires. Therefore,

the focus should be placed on preventative actions.

Transforming forests into fire-resistant landscapes via fuel-treatment activities has been demonstrated to be an efficient strategy in simulated environments for decreasing the expected size and rate of spread of future fires. The need for an effective and practical integrated framework to evaluate these decisions is an open challenge. In this Chapter, we analyzed the outputs from over sixty thousand simulations to demonstrate the theoretical potential of a new method for assessing fire likelihood and its integration into an open-source decision support system. This framework allowed us to systematically design optimized fuel-treatment plans, effectively decreasing the potential of large wildfires by identifying and disrupting the most dangerous fire propagation patterns. Besides its potential applications for landscape managers working with fuel treatment prioritization programs, it could be useful for first responders during firefighting, given the accuracy and execution speed of the model. Our framework can be implemented in any region given an appropriate fire spread model, potentially providing a valuable decision-making tool for researchers to analyze the impact and potential benefits of multiple preventative strategies.

In the following Chapter, we will extend our initial framework to incorporate and condense multiple objectives in the decision-making process. This, with the aim of analyzing the inherent trade-offs faced by the decision makers when performing landscape management activities in real challenging applications. In practice, each action altering the structure of the landscape could impact, positively or negatively, different components of the land such as the economic value of the terrain, environmental impact metrics like the total carbon sequestration of the forest, the accessibility of the area, or even the economic viability of accessing areas for treatment purposes, among several others.

Chapter 4

Decision making under wildfire uncertainty: a multi-criteria extension

4.1 Introduction

In the previous Chapter, simplified landscape management policies were obtained by valuing the landscape regions using a unique objective (e.g., minimize the average expected area burned) or a predefined objective function. However, such an assumption is a simplification of the real challenge faced by decision-makers as multiple parts of the landscape have different values based on factors such as the presence of human settlements and infrastructure, availability of environmental services, and forest health, among others. In this Chapter, we expand these previous attempts by providing an integrated framework to naturally include and weight multiple objectives into the optimization model and analyze the trade-off between present objectives and future protection against wildfire risk. We study three key regions based on their recent fire history, landscape diversity, and demographic variety to quantify the impact of multiple objectives in landscape management. We obtain treatment plans using various combinations of these layers reflecting how different priorities of the decision-makers could affect treatment policies.

Evidence from multiple research areas suggests that the dehydrating effects of climate change have caused a worldwide surge in the number and intensity of fires in the last decade and these numbers are still soaring [1, 121, 168, 290, 363]. These observations are further highlighted when we note that some of the largest and devastating fires leading to heavy human, financial, and infrastructural losses have occurred in the last decade. A closer look at the state of California reveals that despite a decrease in the total number of fires within the state, the total area burned by these fires has increased. In the last ten years, California has been plagued by its most disastrous wildfires [323]. The 2018 Camp Fire is California's most destructive fire recorded, where a single fire destroyed more structures than any other in modern history. This fire is also the most expensive natural disaster in the world in 2018 in terms of insured losses resulting in the loss of 13,696 and the death of over 88 people [38, 110,

131]. Following in level of destructiveness is the 2017 Tubbs Fire in the Napa and Sonoma counties, which destroyed thousands of structures leading to 46 human fatalities [133, 242]. In addition, human efforts in fire fighting are at maximum capacity. These factors suggest that more preventative policies and measures must be taken beforehand to reduce the risk of fire occurrence [102] at the Wild-land urban interface (WUI) and in large areas of forest [73] to protect human lives and maintain an adequate coexistence with nature. In [159], researchers propose a paradigm that considers opportunities in three main dimensions: i) decrease of the potential fire behavior of the landscape, ii) reduction of the potential impact of fire ignition, decreasing the expected losses and number of escape wildfires in fire-prone areas, and iii) increase the capability of fire suppression. As we have already discussed in the previous Chapters, the term Fire-Smart Forest Management (FSFM) has emerged and includes the above concepts.

As discussed in the previous Chapter, the strategic implementation of fuel treatment plans can alter and modify fire behavior and aid in suppression efforts [112, 116, 159, 249]. These fuel treatments consist of actions and procedures such as cutting and clearing wood, prescribed burns, commercial harvesting, and thinning, that can promote fire hazard reduction [3]. In [296], the authors discuss strategies to define treatments using burning probability maps, the area burned, or the flame length. Other research in this area focuses on finding the optimal spatial allocation of prescribed burning activities [11, 218], and designing fire breaks to control fire spreading [292].

The decision-makers that carry out these fuel treatments face questions about how to make such decisions. However, the problem of optimal fuel allocation is challenging due to various sources of uncertainties. In [87], the authors discuss challenges surrounding fuel treatment methods, timing, and the high uncertainty levels in climate and ignition areas over multiple time periods. As a result, deterministic fire simulators such as FARSITE, Prometheus, and Wildfire Analyst [112, 273, 338], which can reproduce fires with and without treatment activities are popular in practice. The main critique to using fuel treatment methods is that due to the difficulty in predicting wildfire occurrence and propagation, the allocation of these treatments generally do not match areas in which future fires occur, leading to wasted investments coupled with economic, human, and environmental losses [39].

In order to address the mismatch between fuel treatment allocation and fire occurrence, we develop a framework that integrates fire spread, optimization, and simulation models (see Chapter 3 and [252]). The study highlights an adaptable metric known as the Downstream Protection Value (DPV), that ranks the impact of treating a unit of the landscape, by modeling a forest as a network and the fire propagation as a tree graph. As discussed in the previous Chapters, the framework requires weather and topography inputs of the forest to run and can be modified by users to incorporate region-specific forest data so as to provide more effective and targeted treatments. Results from the first version of the model value the equivalent volumes of subsets of the forest equally. However, such an assumption is limited in real settings. Different parcels of a landscape may have different values based on multiple factors such as the existence of animal migration corridors, the amount of biodiversity hosted in that region, the presence of human settlements and infrastructure, or the amount of carbon

sequestered.

The main contribution of this study is extending the analysis on the integrated framework proposed in Chapter 3 and [252] to aid decision making under wildfire uncertainty by evaluating the sensitivity of the objective function to key environmental and economic factors. Three key regions in California (Napa Valley, Paradise, and Getty center) are chosen for analysis based on the existence of documented catastrophic events in the last five years and significant variance in vegetation types and demographic variables. We compare unweighted treatment plans against versions including environmental and demographic factors such as carbon sequestration, canopy height and density, population density, and accessibility of the area, as well as expected future fire behavior. Using multiple combinations of these layers, we generate different treatment plans that reflect how different priorities of the decision-makers could affect the treatment policies. We then analyze the trade-off between maximizing the decision-maker utility function and protecting the land against future expected losses due to wildfire with the aim of finding robust treatment plans.

4.2 Data Extraction and Processing

The California region with its established seasonal fires is the focus region in our study because of the existence of documented and destructive fire events. The default framework relies on weather and topographic data obtained from various sources. In addition, we discuss how to generate the main layers that serve as weights for estimating the risk associated with each unit of landscape in our optimization module. We group the layers extracted for the objective function into Environmental factors (e.g., canopy density/height) and Demographic characteristics (e.g., population density, accessibility). Additionally, we split the data into a training and test set. During training, we fine-tune key parameters in our model using the training data set. The test set allows us to measure the performance of our model by assessing how well our model results compare to real data.

The data consists of a combination of shapefiles and GeoTIFF files. For the most part, these layers are already aligned and have the same resolution (30 x 30 m). We transform the original data into a series of rasters to be ingested into our framework. Data enters into the framework using two main approaches: (1) local data provided by the user and (2) semi-automatic collection and processing of online assets (e.g., population density, climatic conditions that can play a fundamental role in the propagation of the fire) available in Google Earth Engine (GEE) [144], to easily generate a consolidated dataset (Figure 4.1). This latter method relies on Python scripts and can be used as an independent module for performing any kind of query in GEE. Finally, the data is automatically processed and formatted to be ingested into the different modules of the proposed framework.

Environmental factors

The building blocks for the data layers in this category are obtained from the Landscape Fire and Resource Management Planning Tools (LANDFIRE) [285] open repository. This data hub consists of a shared program involving the participation of the U.S. Department of Agriculture Forest Service and the U.S. Department of the Interior, providing support for fire and vegetation simulators created and used by the US Forest Service. The Canopy bulk density (CBD) and the Canopy height (CH), each a grid with resolution (30x30 m) are obtained directly from the database. CBD is a measure of the density of the landscape canopy which is the portion of vegetation above ground. Canopy height as implied by its name measures the height in meters of the landscape canopy.

Forest. We combine the CBD and CH layers in various manners to obtain the different layers constituting the environmental factors category including estimating the area of green vegetation in the forest (forest health) and carbon sequestration volume. The CBD and CH layers are averaged together to form a new layer that serves as a proxy for forest health. The CBD layer indicating vegetation density and the CH measuring vegetation height provide insights about the health of the landscape vegetation.

Carbon Sequestration. The carbon sequestration (CS) layer is estimated following the simple method proposed in [128] for each unit in our landscape grid as follows:

1. Estimate the total (green) weight of the trees. Using the canopy density, we obtain the canopy area by multiplying the density by the area of a unit cell. Then the canopy area is multiplied by the canopy height to obtain a volume. Once we have the volume of green in each cell, we use Huber's formula [350] to estimate the weight using the equation:

$$\text{Weight} = \text{volume} \times \text{green density}$$

where green density represents the total density of wood and bark combined (CBD).

2. According to previous research [106], the average dry weight for different temperate tree species in the United States is about 72.5% of regular weight. We use these results to approximate the average dry weight of green for each cell.
3. Next, we estimate the weight of carbon using findings from [47], which state that the carbon weight is about 50% of regular tree weight.
4. Finally, we calculate the average weight of carbon dioxide sequestered in the tree per cell using:
 - a) The atomic weight of CO_2 is $\text{C} + 2 \times \text{O} = 43.999915$.
 - b) The ratio of CO_2 to C is $43.999915/12.001115 = 3.6663$

Therefore, to determine the weight of carbon dioxide sequestered in the tree, we multiply the weight of carbon in the tree by a 3.6663 factor.

Fuel vegetation type. The fuel layer With a 30 by 30 m resolution that we require for our simulations is obtained from the LANDFIRE [285] publicly available repository¹. It provides a categorical grid at a national level representing the forty Scott and Burgan fire behavior fuel models lastly updated in 2014. For a comprehensive analysis and description of all fuel types, their characteristics, and experimental parameters, see [307]. State-level data (California) is locally extracted and uploaded into GEE to consolidate it with the additional layers of the study. This is the main input for calculating the fire rate of spread in the simulation model.

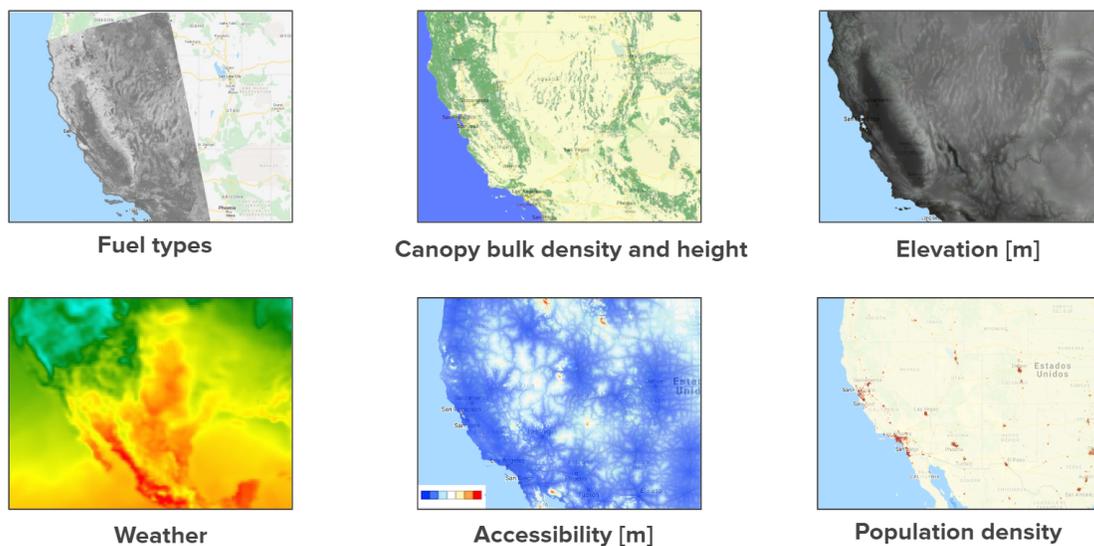


Figure 4.1: Example of layers included visualized in GEE. All layers are obtained for California and consolidated into a 30 by 30 m. multi-band raster.

Weather and moisture scenarios. The rate of Spread (ROS) mainly depends on the type of forest fuel that goes into combustion and the wind speed. However, another influencing factor is the moisture content of dead and live vegetation present in the forest. This variable is crucial for estimating the surface ROS as certain land-covers (e.g., grass types) tend to significantly modify their propagation patterns depending on their humidity levels. In order to capture and simulate interesting situations for practitioners and researchers, multiple value thresholds are empirically studied. According to Scott & Burgan [307], these

¹https://www.landfire.gov/bulk/downloadfile.php?FNAME=US_140_mosaic-US_140FBFM40_20180618.zip&TYPE=landfire

interesting humidity levels are represented by four scenarios: D1L1, D2L2, D3L3, and D4L4, from the driest (D1L1), where the fire tends to propagate faster, to the wettest (D4L4), with the opposite effect. Weather scenario files that describe the evolution of the temperature, wind speed, and wind direction are obtained from the historical time series of the closest weather station (with respect to the centroid of each instance) available for the simulated fire duration. Each fire is simulated for 12 hours under the D1L1 scenario, to be able to capture relevant propagation patterns.

Demographics. The estimated population densities (number of persons per square kilometer) for the years 2000, 2005, 2010, 2015, and 2020 are extracted from the Gridded Population of World Version 4 (GPWv4), Revision 11 dataset² at a resolution of 30 arc-second grid cell and averaged. Accessibility to cities, which measures the land-based travel time (minutes) to the nearest densely-populated areas with 1,500 or more inhabitants per square kilometer is obtained from the Malaria Atlas Project³ at a 30 second-arc resolution for 2015.

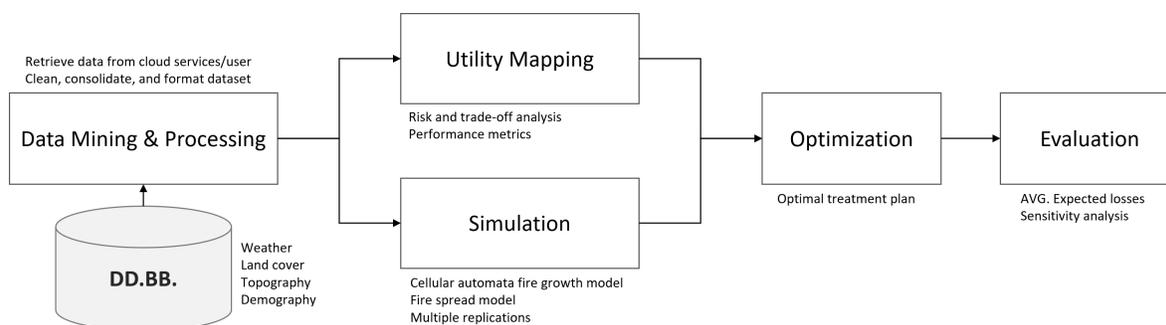


Figure 4.2: Framework schematic. Data is retrieved from cloud services and local user inputs. Decision-makers define relevant objectives by analyzing the trade-off between multiple variables. Once the data is processed, multiple simulations are performed to estimate the impact of future wildfires in the landscape. An optimization model is fed with the outputs from the utility mapping and simulation models. Finally, results are evaluated by estimating the average expected losses due to wildfire, as well as analyzing the sensitivity of the optimal treatment plan.

4.3 Wildfire simulation

Following the developments of the previous Chapters (see section 3.2, Chapter 3 for details), a region of interest is modeled as a two-dimensional lattice with an underlying network

²<https://sedac.ciesin.columbia.edu/data/collection/gpw-v4>

³https://malariaatlas.org/research-project/accessibility_to_cities/

structure to represent the connectivity between cells/nodes in the grid. Cells represent a homogeneous area with similar characteristics, focusing on points of interest such as potential propagation sources, densely populated areas, natural reservoirs, among others. Each cell from the input data corresponds to a node in an undirected grid graph $G = (\mathcal{N}, \mathcal{E})$. Nodes $i \in \mathcal{N}$ is associated with relevant attributes (e.g., population density, fuel type) included as part of the risk evaluation and further optimization models. Nodes are connected by edges $e \in \mathcal{E}$ with weights (the distance between cells) generated by the structure of the region. In the current version of the study, each cell has at most 8 neighbors to simplify the analysis of the framework. However, the model can be easily extended to a general case (graph) where nodes can be connected with any other in the lattice as long as they share information or are related by relevant variables for the model.

Adapting the simulation framework discussed in Chapters 1 and 2, the fire spreads following a messaging process between the cells of G . The intensity of these signals is represented by the rate of spread (ROS) obtained from an independent fire spread model (BEHAVE) integrated into the simulator [69] that models the fire behavior for static conditions in the U.S. based on empirical studies. It is used to update the fire progress between neighboring cells at every time-step t . From here, a directed tree graph – denoted as *shortest-path tree* – is obtained, with a root at the ignition node i , $\mathcal{T}_i = (\mathcal{N}_i, \mathcal{E}_i)$ where $\mathcal{N}_i \subseteq V$ contains the burned cells and \mathcal{E}_i the directed edges representing the propagation trajectories of the fire within the region (see Chapter 3).

Simulator parameters may need calibration in order to reproduce realistic propagation patterns or capture the impact of previously unseen conditions. Multiple approximations during the implementation of the spread models and inherent noise within the data may inaccurately represent the expansion of the fire. To account for this situation, we apply the automatic adjustment following the work in Chapter 2 of the main parameters of the model via a hybrid AI-Optimization procedure that aims to minimize the differences between the simulated and historical fire scars of the region of interest. Using this approach, we can automatically adjust the ROS estimations to accurately represent observed fire perimeters. Therefore, decision-makers are able to automatically adjust the fire spread model to account for variations in the fire behavior and conditions of the area, obtaining more accurate simulation results.

4.4 Risk analysis

Relevant features are mapped onto a common scale and weighted to account for their relative importance for the decision-maker. In order to condense their information into a single utility function, each feature is mapped unto a common scale $[0, 1]$ using an adequate function (e.g., linear) according to the expected impact of the feature in the landscape. Thus creating a single matrix representing the original $n \times m$ grid containing the value for each node in landscape ($NV \in \mathbb{R}^{n \times m}$), different convex combinations of the relevant values-at-risk (raster layers) are utilized. Let $\mu_k \in [0, 1] \forall k \in K, \sum_{k \in K} \mu_k = 1$ with K being the total number of

layers included in the study, we thus combine the layers as follows:

$$NV = \sum_{k \in K} \mu_k L_k \quad (4.1)$$

where $L_k \in \mathbb{R}^{n \times m}$ is the matrix of dimensions $n \times m$ containing the grid values of layer k . Following this framework, we obtain a consistent risk function across the entire graph. The detailed procedure is as follows:

1. Given a set of K features representing characteristics of the nodes, we map them into a common scale between $[0, 1]$. The mapping function (e.g., linear) from raw feature values to the $[0, 1]$ interval is selected by the researcher according to the impact of each feature in the construction of a global utility/cost function. For example, if the protection of nodes with higher population density is prioritized, an increasing non-linear function can be applied, where densely-populated nodes are associated with values near to one, while sparsely populated areas are mapped to near-zero values (see Figure 4.3).
2. This procedure is repeated for all K features, obtaining a set of normalized variables.
3. Correlated and complementary variables are combined into meaningful categories by weighting individual features with weights. As an example, canopy bulk density and canopy height could be summarized into a *Forest* category that equally weights both variables. Similarly, accessibility and population density layers could be condensed into a *Demography* category.
4. Once all categories are generated and normalized, a global utility function is calculated repeating the weighting procedure.
5. Gaussian kernels are applied to smooth the distribution over the landscape/grid. This avoids abrupt changes in the utility function as well as accounts for the intrinsic correlation of the cells in the landscape dynamics (Figure 4.6).

Using this framework, decision-makers are able to condense any number of features into unique values associated with each node of the graph (NV), as well as generate a series of scenarios to evaluate the impact of certain features and their weights. In this way, each node of the grid is characterized by a unique value incorporating all variables. Using different weight combinations allow us to control which factor we want to weigh more in each instance (see Figure 4.5). Next, we utilize our proposed risk model, downstream-protection-value (DPV , [252]), which aims to measure the value of what is affected downstream in the network given the risk of a certain node. As exposed in the previous Chapter, this model captures the role of a node in propagating the risk through the landscape based on the connectivity and relative influence of different nodes in the entire system. In the extended

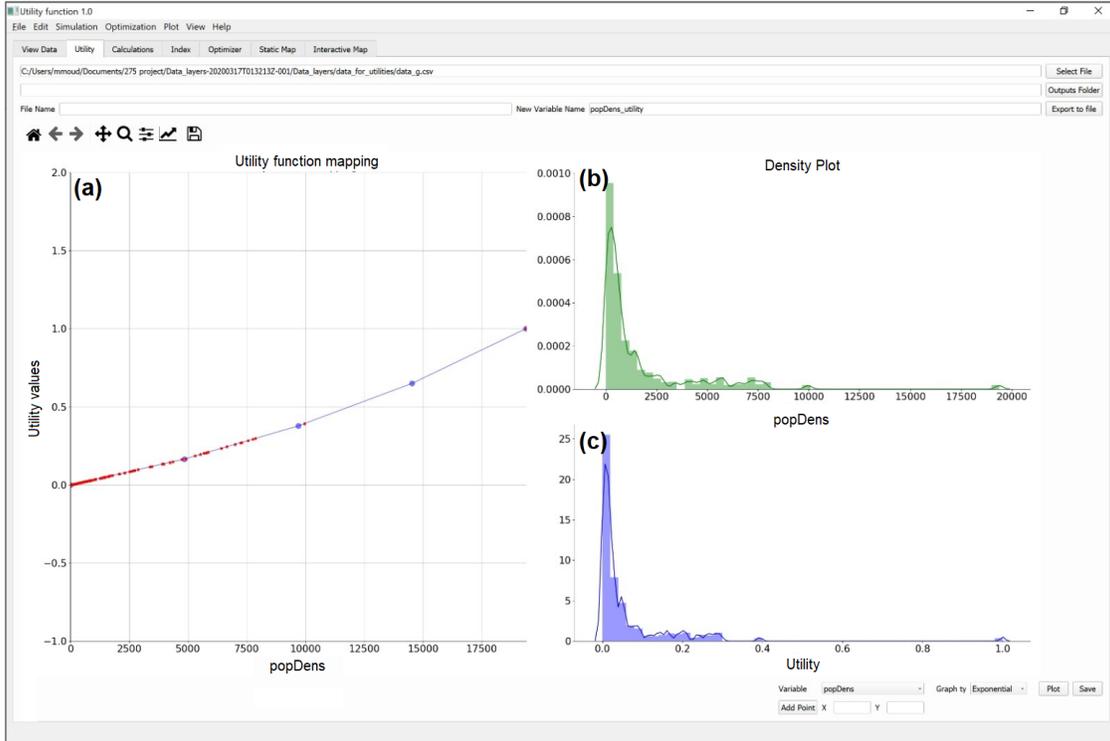


Figure 4.3: Utility mapper application. (a) The original values of the population density layer (x-axis, popDens) are mapped to the $[0,1]$ interval following an exponential function (y-axis). (b) Density plot for the original variable values. (c) Distribution of the mapped $[0,1]$ feature (called utility) following the applied transformation.

version of the model, we define the DPV of node i inside the landscape network as:

$$DPV(i) = \alpha_i \sum_{j \in \mathcal{N}_i} NV_j \quad (4.2)$$

where α_i is a weight factor for node i , e.g., the number of connections inside the region, allowing the prioritization and management of different zones; \mathcal{N}_i is the set of nodes of the network that are affected by fire propagation from node i , and NV_j the value of node j incorporating all the relevant variables included to capture the potential losses caused by wildfire. Researchers can represent and evaluate multiple scenarios by providing different weights to the components of the risk function or to the final node value, thus obtaining variations of DPV . For example, NV values of nodes playing a fundamental role in the connectivity of the network could be weighted, among several options, by the degree of the node to highlight its importance in the propagation of fire. In addition, thanks to the nodes' additive property, there exists a natural extension from nodes to larger units (e.g., stands). This

goes in hand with a practical implementation since authorities take decisions over certain areas with common characteristics instead of individual units, satisfying operational/logistic constraints.

An optimized resource allocation plan

Formulating an extended version of the explicit optimization model presented in Chapter 3, Section 3.2, we solve the problem of finding the connected cells that maximize the total utility/protection value considered under wildfire risk, subject to specific constraints provided by the decision-maker (e.g., budget). As presented in our previous studies located in Canada, the connectivity constraints are imposed to mimic realistic scenarios where the protection of the landscape must be performed within connected patches to account for operational constraints. If needed, these constraints can easily be relaxed by the decision-makers, running a simplified version of the optimization model, or even implement their own algorithms. In this way, we obtain an optimized treatment plan that identifies the set of units that, once protected, will significantly disrupt the fire propagation while taking into account the impact on relevant features of each node via the provided NV values. We introduce $\lambda \in [0, 1]$ as the trade-off factor between the utility function defined by the NV matrix and the DPV heatmap obtained after simulating R wildfires in the area. We define the objective function as the convex combination $U(\lambda) := \lambda DPV + (1 - \lambda)NV$, where larger values of λ indicate that wildfire risk becomes more relevant to the decision-maker when allocating resources through the network ($\lambda \rightarrow 1$) while smaller values represent the case when the NV values play a more relevant role for prioritizing the treatments ($\lambda \rightarrow 0$). Therefore, different optimized resource allocation plans are obtained depending on the objective and risk assessment of the decision-maker represented by the λ parameter, balancing the maximization of the utility function and the protection of the region against the future propagation events (See Appendix D, Figure D.1).

Mathematical Formulation

Following the model presented in Chapter 3, Section 3.2 [97, 252], the $\bar{\lambda}$ -connected version of the Protection Value Problem ($PVP_{\bar{\lambda}}$) can be formulated as the following Mixed-integer programming model:

$$(PVP_{\bar{\lambda}}) \quad \max \sum_{i \in V} U_i(\bar{\lambda}) = \sum_{i \in V} (\bar{\lambda} DPV(i) + (1 - \bar{\lambda}) NV_i) x_i \quad (4.3)$$

$$s.t. \quad (3.3) - (3.14) \quad (4.4)$$

Therefore, we modify the original objective function to incorporate the trade-off between DPV (i.e., weighted fire protection) and the inherent utility associated with the node represented by NV . The importance of each objective is then modeled with different λ values, providing a set of comparable solutions for the decision-makers according to their objectives.

As discussed in Chapter 3, this is an NP-hard problem, being a variant of the MWCSP [166]. Using the two-step solution method described in Chapter 3, Section 3.2 the problem is able to converge faster because we start the second step (i.e. solving the MIP formulation with CPLEX v12.9) with a high-quality feasible solution. This reduces the searching space of the method significantly (upper and lower bounds), improving the convergence and memory usage of the optimization algorithm. Thus, we obtain optimal solutions in reasonable (less than 1 minute on average) solving times in all our experiments. We note that our solutions are guaranteed to be globally optimal because we obtain 0% optimality gap solutions from the exact MIP formulation. Although multiple solutions can exist, these are extremely rare in practice.

4.5 Case study areas

The areas in California selected for the case study are chosen based on factors including significant variation across fuel types, the existence of documented catastrophic fires in the last five years, and the presence of strongly influencing demographic factors. The three areas selected for this study are the Getty center, Napa Valley, and Paradise. Figure 4.4 denotes the terrain of the three case studies areas. In addition, a summary of the key features of our three study instances including area, mean elevation, dominating fuel types, and elevation range are presented in Table 4.1.

Napa Valley. Napa County, also known as Napa Valley is recognized worldwide as a premium wine region. In addition, the region is also responsible for the production of many agricultural crops. According to the U.S. Census Bureau, the county has a total land area of about 748 square miles and a population of 137,744 as of 2019 [80]. However, this region has also suffered a number of destructive fire events. A recent article mapping all fires in the region from 1950 to 2019 shows that the fires have been getting larger and more destructive [192]. Most notably the 2017 “Tubbs” fire in the Napa and Sonoma counties is the second most destructive fire recorded in California’s history [162]. The long fire history experienced by this region coupled with its landscape being suitable for the farming of different agricultural products particularly grapes from vineyards makes it a very relevant and interesting area to study in California. The total number of nodes and edges conforming this instance are 9,309 and 72,272, respectively.

Paradise. Paradise is a small town located in Butte County in California that has experienced the most destructive fire in California’s history at the end of 2018. In 2018, the population of Paradise was about 26,800 with a land area of about 18.31 square miles [80, 162]. However, the population numbers after the fire event are unknown as over 9300 were displaced and relocated during and after the fire [195]. We use population numbers before the fire as an estimate for our analysis. This town is a unique region to study because of its

high fire risk, limited accessibility to, and relative isolation from neighboring towns. With a total number of 11,477 nodes and 81,581 edges, it is the largest instance of the study.

Getty center. The Getty center, a campus of the Getty Museum and its surrounding regions located in Los Angeles is selected due to its 2019 fire which forced the museum to temporarily close its doors [77]. This region is different from the other regions of interest because it is located in an area with a more built-up environment and less vegetation. We speculate that the demographic factors will have a stronger influence on our analysis of this region. Modeling the instance as a network, it consists of 5,454 nodes connected by 41,166 edges.

Table 4.1: Summary of instances’ main characteristics. For each instance, we provide the total area in hectares, the average elevation and its range in meters, the dominant flammable fuel of the terrain following the fuel type layer characterization, the total number of different fuels available in the region, and the number of edges conforming the network used for the optimization model connecting the flammable cells (in any direction).

Instance	Area [ha.]	Mean elev. [m]	Elev. range [m]	Dominant flammable fuel	# Fuel types	# Edges
Napa Valley	9,540	376.74	[131, 724]	SH2 is woody shrubs and shrub litter	16	72,272
Getty center	11,102	221.39	[39, 596]	GS2 is grass and shrubs combined	11	41,166
Paradise	13,433	449.71	[0, 740]	TL6 is moderate load broadleaf litter	17	81,581



Figure 4.4: Land cover representations. The three case study areas are depicted with a hill shade effect where different colors represent the fuel types characterizing the instances following the Scott & Burgan [307] classification system.

Utility Mapping. For each of the three instances we study, we utilize four layers in our utility function including forest, volume of carbon sequestered, accessibility, and population

density layers. We select mapping functions to highlight the priorities of the decision-maker for every layer. These functions serve to map the values of our current layers to a range between 0 and 1 using a pre-specified distribution. We Choose an exponential mapping function for the population density layer because we want an increasingly high utility as the number of people in a region increases. In addition, we use the inverse function to map our accessibility values to the range between 0 and 1. This function is selected to reflect increased utility for regions that are closest to densely populated areas. Finally, a linear function is used to map the volume of carbon sequestered and forest layers to the appropriate ranges. Our framework allows provides the option to change the distribution of the mapping functions and customize these functions based on the decision-maker’s most important considerations. In order to combine the different features into a single utility, we select five convex combinations of our different layers for our experiments. First, we use an average combination where all four layers are weighted equally. Then, we create four feature dominant layers, where the selected dominant feature contributes 75% to the final utility whereas the remaining three layers equally split the remaining. These five utility combinations are created for each of our three instances (Napa Valley, Paradise, and Getty center).

4.6 Experiments

We model five different NV values for each instance by modifying the μ weights vector associated with the different four categories following the procedure described in Section 4.4. In this way, we obtain a balanced weight function $NV_{equal} = \sum_{k \in K} 0.25L_k$ where all layers are weighted identically and four variations where a dominant layer is weighted by $\mu_{dom} = 0.7$ and $\mu_j = 0.1$ for $j \neq dom \in K$ obtaining NV_{forest} , NV_{access} , NV_{carbon} , and $NV_{population}$.

A total of $R = 100$ replications with random ignitions and defined D1L1 weather scenarios are performed in the simulation module to obtain the final DPV heatmaps using the previously generated NV values as the node weights, multiplied by $\alpha_i =$ number of neighbors connected to node i . For each instance and NV value combinations, we generate utilities $U(\lambda)$ with $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$. Each combination is then solved in the optimization module for all treatment levels $tf \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.5\}$, solving a total of 450 optimization problems. Optimal solutions $X^*(\lambda, tf)$ indicating the selected cells are recorded for each combination.

Finally, we evaluate the average expected losses, $\mathbb{E}[Losses(X^*(\lambda, tf))]$, due to future wildfire events in each instance, given the output of the resource allocation plan. For this, we estimate the expected damage provoked by future fires – discounted by a $\gamma \in (0, 1]$ factor set to $\gamma = 0.9$ – using $R = 100$ simulations in a modified landscape where the selected cells from the optimal solution of the $PVP_\lambda(tf)$ model are transformed into non-flammable ones. We use this as a simplified version of a more realistic setting where certain fuel treatment actions (or resource allocation) will decrease the fire susceptibility (or effective ROS) of a certain cell but it will be still flammable. The pseudo-code summarizing all the experimental

steps can be found in Algorithm 1.

Algorithm 1 Experiments Pseudo-code

```

1: procedure EXPERIMENTS
2:   for  $dom \in \{forest, access, carbon, population, equal\}$  do
3:     if  $dom == equal$  then
4:       Calculate  $NV_{dom} = 0.7L_{dom} + \sum_{k \in K: k \neq dom} 0.1L_k$ 
5:     else
6:       Calculate  $NV_{dom} = 0.25 \sum_{k \in K} L_k$ 
7:       Calculate  $DPV(i) = \alpha_i \sum_{j \in T_i} NV(j)$ ,  $\forall i \in V$ 
8:       for  $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$  do
9:         Calculate utility  $U_i(\lambda) = \lambda DPV(i) + (1 - \lambda)NV(i)$ ,  $\forall i \in V$ 
10:        for  $tf \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.5\}$  do
11:          Get optimal solution  $X^*(\lambda, tf) \leftarrow \text{Solve}(PVP_\lambda(tf))$ 
12:          Calculate  $\mathbb{E}[Losses(X^*(\lambda, tf))] \leftarrow \text{Simulation}(X^*(\lambda, tf), R = 100)$ 
13:          Calculate  $\Delta_{tf}U(X^*(\lambda, tf)) := U(X^*(\lambda, tf)) - \gamma \mathbb{E}[Losses(X^*(\lambda, tf))]$ 

```

Computational implementation

As described in Chapter 1, the Cell2Fire fire-growth simulator is implemented in C++ using the boost and omp libraries [300, 333] to allow shared memory parallel execution. The decision support system which wraps the simulation module and processes all relevant outputs is programmed in Python. Statistics and visualizations are processed using the known Pandas, Numpy, and Seaborn libraries. Network structures are managed with the networkx package [150], generating outputs such as shortest-path trees and providing the users a variety of complex network metrics out-of-the-box such as betweenness centrality for developing their own fire risk indexes. we implement derivative-free optimization algorithms [95] in our framework using the NLOPT package [167] (Chapter 2). Meta-Heuristics (e.g., genetic algorithms) follow the DEAP [127] library framework. In addition, we rely on the PYOMO modeling language [154], a flexible package to generate linear/non-linear models and solve them via an open-source or commercial solver depending on the user needs to embed mathematical programming models in the framework’s optimization module. The utility mapper standalone application and scripts are programmed in Python using the PyQt5 package and compiled using the Pyinstaller package.

Experiments are performed in a daily use laptop with I7-4200 2.1 GHz processor, RAM Memory 8 GB DDR3, and Ubuntu 14.0 OS. All codes are available for public use⁴.

⁴http://www.github.com/cpaismz89/DPV_Utility

4.7 Results and Discussion

The utility functions in our experiments are obtained by combining raster layers as described in Section 4.4 and the different utility combinations are explained in 4.6. We conceive these different combinations to mimic decision-makers having multiple objectives, where one of the goals is more important than the remaining ones in the decision-making process. In Figure 4.5, we present the utility maps for all proposed combinations of the primary layers for the three instances we explored. We observe that the combinations produce very different heatmaps which we hypothesize will be translated into different optimal treatment decisions.

Utility and wildfire risk trade-off

DPV values are obtained for all the generated utility heatmaps following Eq. (4.2) with NV_j representing the value of cell $j \in V$ from the calculated utility layer. Looking at the DPV matrices (Figure 4.6), we clearly observe the most likely wildfire propagation patterns after performing $R = 100$ replications for each instance (first column) where lighter colors highlight those sections of the landscape with higher DPV, i.e., the nodes that play a fundamental role in propagating the fire between different areas of the land. As expected, the DPV matrix obtained for the Napa Valley instance covers a significant portion of the landscape (88.63%) since it is the one with the largest proportion of flammable fuel types, representing 97.5% of its total composition. On the other hand, Paradise and Getty center instances include a significant proportion of non-flammable nodes representing urban settlements/rocky areas (14.56% and 50.87%, respectively) leading to more focused DPV heatmaps surrounding those areas, covering 40.3% and 37.27% of their total size, respectively.

The expected area burned and expected utility losses under the current conditions, i.e., when no resource allocation plan is implemented, can be seen in Table 4.2. From the experiments, we observe a significant impact on the Napa instance with an expected area burned of 2,055 ha. representing 22% of the landscape. This is translated into expected losses close to 20% among all utility functions with respect to the total value available. In the case of Getty center, we expect a 19% of the area burned due to future wildfire events, with an impact on the utility functions varying from 14.32% (Forest dominant utility) to 16.18% (Access dominant utility). Finally, due to the characteristics of the Paradise instance – the distribution of the non-flammable fuel types – a 7% of the total area of the landscape is expected to be affected by future wildfire events under the tested conditions, leading to expected losses between 3.91% (Population layer) to 7.58% (Forest layer).

Analyzing the most common fuel types involved in the propagation patterns identified when calculating the DPV, we observe grass and shrubs combined (GS2, 36.84%), grass though small amounts of fine dead fuel (GR2, 27.18%), and woody shrubs and shrub litter (SH7, 20.40%) for Napa Valley; woody shrubs and shrub litter (SH7, 50.73%) and grass and shrubs combined (GS2, 46.14%) in the area near Getty center; and woody shrubs and shrub litter (SH7-SH5, 62.49%) and grass and shrubs combined (GS2, 23.14%) for the Paradise

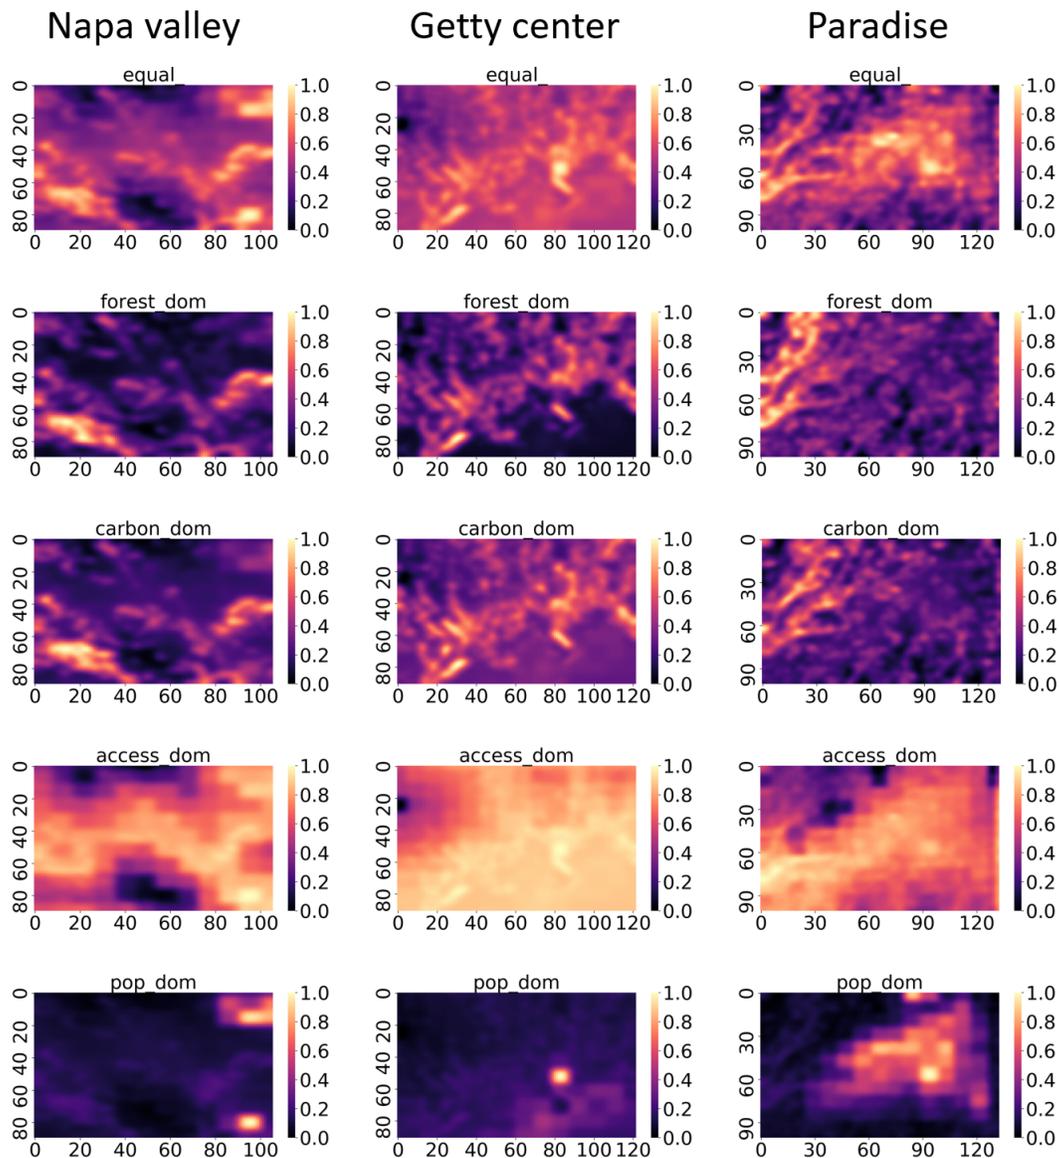


Figure 4.5: Utility heatmaps for all proposed convex combinations of the main four categories for each instance (columns). The first row represents a balanced combination of all four categories ($\mu_i = 0.25 \forall i$). The suffix *dom* indicates that the dominant category was weighted by $\mu_{dom} = 0.7$ and the remaining three categories with $\mu_j = 0.1, \forall j \neq dom$.

instance. This information allows the decision-maker to gather relevant insights about the most dangerous sections of the landscape in terms of wildfire risk and identify the set of potential actions to mitigate their impact when implementing the solution obtained from the optimal resource allocation plan.

Table 4.2: The expected area burned and expected losses for all utility functions as a percentage of the total instance area and the total utility available (heatmaps) per instance, respectively. Expected values are calculated from $R = 100$ independent wildfire replications, weighting all simulations equally, and without any intervention of the landscape.

Instance	$\mathbb{E}[Burned]$ %	$\mathbb{E}[L(U_{Equal})]$ %	$\mathbb{E}[L(U_{Forest})]$ %	$\mathbb{E}[L(U_{Carbon})]$ %	$\mathbb{E}[L(U_{Pop})]$ %	$\mathbb{E}[L(U_{Access})]$ %
Napa Valley	22.08%	20.73%	19.66%	20.06%	18.95%	21.75%
Getty center	18.97%	15.27%	14.32%	14.66%	14.43%	16.18%
Paradise	6.94%	6.33%	7.58%	6.83%	3.91%	6.94%

Resource allocation plans

As described in Section 4.4, we analyze the trade-off between the expected losses due to future wildfires in the area after applying the optimal resource allocation plan – assuming full protection of the selected nodes for simplicity – and the protected value due to the implementation of this plan by combining the utility layer generated by the decision-maker and the DPV matrix obtained from the simulations via the λ parameter. Depending on the expectations of the decision-maker and his/her level of risk aversion, different λ values should be tested and selected for a particular region and context. In the rest of this section, we will focus our attention on three interesting results where the trade-off between present value and future protection plays a crucial role in the decision-making process. This analysis can be performed for all combinations of utilities, treatment fractions, and instances, providing the decision-maker with a comprehensive set of results (see Tables D.1, D.2, and D.3 in Appendix D) and quantitative support for establishing the optimal point to balance the trade-off between present utility and future protection of the landscape.

To illustrate and analyze the inherent trade-off between present utility and the future protection of the landscape, we observe the results for the Napa Valley instance when identical weights are provided for all categories conforming the utility layer (Figure 4.7-(a)) and a 25% of the landscape is protected. As expected, the present/raw utility tends to decrease as $\lambda \rightarrow 1$ since the decision-maker is sacrificing present utility by focusing more resources in those locations where the fire will likely propagate, prioritizing the disruption of future wildfire events instead of the original utility function. However, present utility values are overestimated when $\lambda \rightarrow 0$ as the decision-maker oversees the impact of future wildfires, not accounting for this risk. From the graph, we can observe that the best performance in terms of discounted utility is attained when $\lambda = 0.5$ (1017.39 ± 186.06), obtaining the best balance between raw utility and expected wildfire losses using a discount factor of $\gamma = 0.9$.

Analyzing the gaps between the present utility and the discounted function, we observe a clear decreasing pattern as $\lambda \rightarrow 1$ in terms of total utility value variations⁵, with differences of 38.23%, 30.6%, 23.65%, 15.75%, and 7.41 %, respectively. This is aligned with our expect-

⁵Calculated as $\mathbb{E}[Losses(X^*(\lambda, tf))] / U^*(\lambda, tf)$

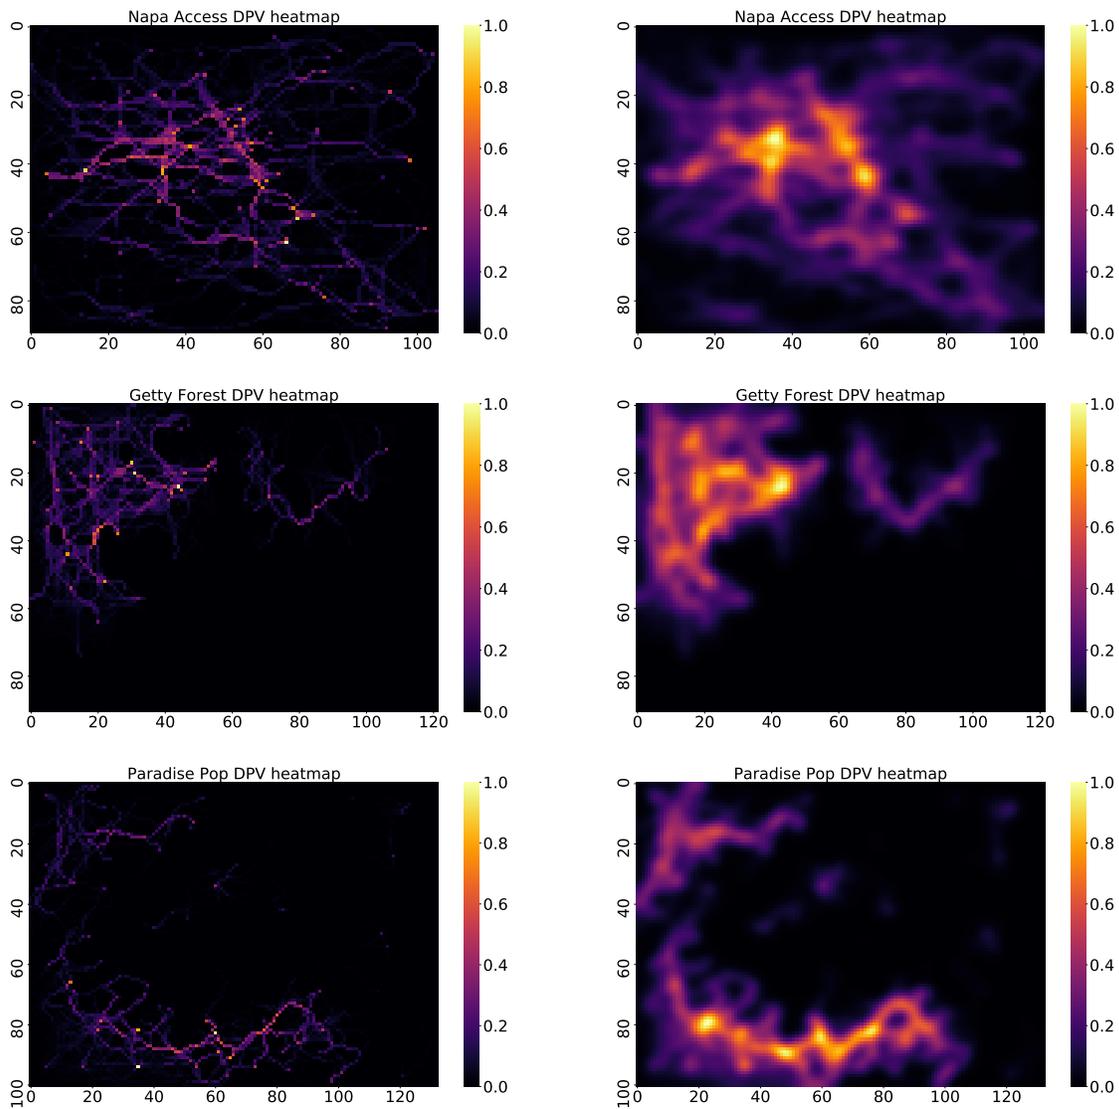


Figure 4.6: Raw (left) and smoothed (using a Gaussian kernel - right) DPV heatmaps calculated for Napa Valley, Getty center, and Paradise instances using the Access, Forest, and Population density layers as the dominant layers for the NV function, respectively. Lighter cells increasingly highlight the nodes playing a fundamental role in propagating the fire to the rest of the landscape.

tations: as λ is increased, the lands are better prepared for future wildfires, minimizing the gap between today's total utility and the discounted function by sacrificing present value. Looking at the distribution of the discounted utility as a function of λ (Figure 4.7-(b)), we

observe that increasing the weight of the expected wildfire risk ($\lambda \rightarrow 1$) results in a more compact distribution of the discounted utility $\Delta U(\lambda)$, as seen in the graph. This is consistent with the fact that higher λ values lead to better protection plans by sacrificing the value of the utility function, a trade-off that the decision-maker will analyze to decide which resource allocation plan is aligned with her expectations and goals.

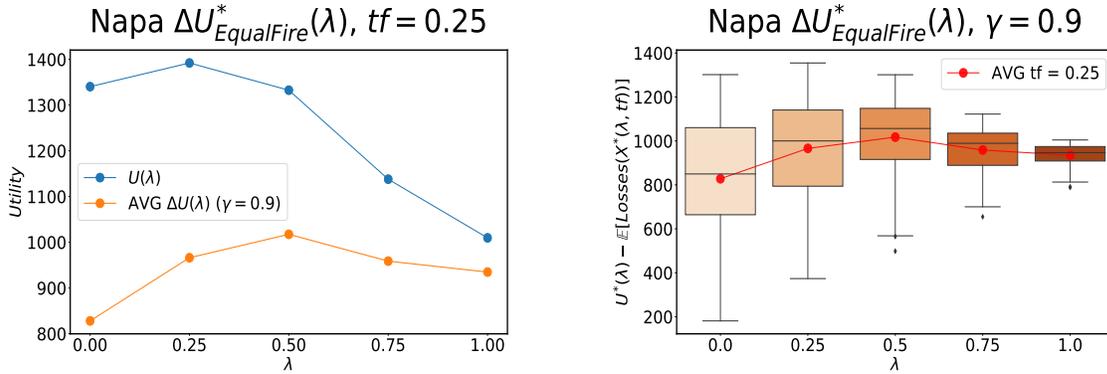


Figure 4.7: (a) Napa Valley instance raw utility (blue) $U(\lambda)$ weighting all categories by identical weights and average discounted utility (orange) $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 25%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 25% of the landscape. Average values are highlighted with red dots.

We then focus our attention on the results obtained for Paradise with a carbon-dominated utility when treating 50% of the total landscape. Visualizing the optimal plans for multiple λ values (Appendix D, Figure D.1), we observe that the resource allocation plan is not as sensitive as in the case of Napa Valley, but we can still observe differences across the different levels of λ to analyze the trade-off between present value and the protection of the landscape. This is mainly associated with three factors: (1) we are treating 50% of the land so there exists a larger overlap between the optimal plan and the propagation patterns identified in the DPV matrix, (2) the distribution of the carbon dominated utility matches the most relevant DPV spread lines, and (3) the instance, similar to the results from Getty center, presents a significant amount of non-flammable nodes, thus limiting the potential fire spread paths.

This is translated into significantly smaller gaps between the present/raw utility value and the discounted one for all λ levels (3.15%, 2.45%, 0.86%, 0.4%, and 0.04%, respectively). As seen in Figure 4.8-(a), both functions converge to an almost identical value for $\lambda \geq 0.5$. This situation indicates that optimal plans giving at least 50% of weight to the DPV layer are able to significantly control and mitigate the future expected losses due to wildfire events. Observing the distribution of the discounted utilities in Figure 4.8-(b), it can be seen that the

value and variation in the discounted utility (y-axis) are significantly affected by λ , obtaining different levels of risk. From the results, the plan obtained when $\lambda = 0.5$ arises as a good solution, balancing the raw utility and expected losses as well as being characterized by a compact distribution (2250.25 ± 17.58). Values of $\lambda < 0.5$ lead to greater average discounted utility (red dots) but incur in more risk and potential negative outcomes (e.g., low discounted utility value outliers in $\lambda = 0.25$). On the other hand, values with $\lambda > 0.5$ sacrifice a portion of present utility to decrease the impact of expected future losses (5.1% and 14.5% w.r.t. $\lambda = 0.5$, respectively).

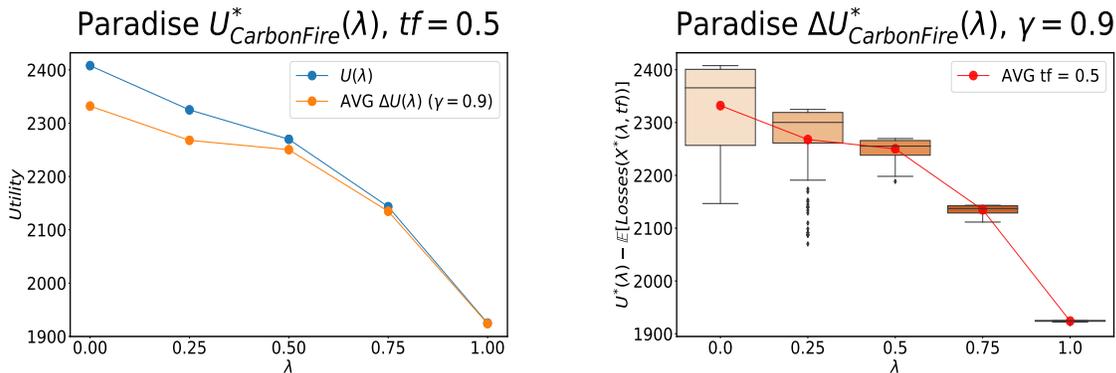


Figure 4.8: (a) Paradise instance raw utility (blue) $U(\lambda)$ with carbon as the dominant category and average discounted utility $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ (orange) including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 50%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 50% of the landscape.

Finally, we analyze the results obtained for the Getty center instance, when the utility function is dominated by the accessibility layer and only 15% of the landscape receives treatment. Contrary to the previous results, we observe an increasing pattern in the discounted utility (Figure 4.9-(a)) as more weight is provided to the DPV matrix, this is, the optimal plan is mainly aligned with the mitigation of future wildfire losses instead of the current benefit ($\lambda \rightarrow 1$). These results indicate that, for this particular experiment, focusing the attention only on the present objective function value for selecting the nodes to be treated/protected is not the most efficient solution as there is no significant intersection between the treatment plan and the most likely propagation patterns experienced in the landscape, leading to larger expected losses as $\lambda \rightarrow 0$. Analyzing the gap between both curves, we observe differences of 148.48%, 119.77%, 74.86%, 49.17%, and 31.64% as we increase the value of λ , respectively. Therefore, $\lambda = 1$ arises as a robust option under the current experimental parameters.

Following the discussion, we observe how the distributions of the discounted utilities (Figure 4.9-(b)) with $\lambda < 0.5$ are particularly wide with a significant bias to the bottom (-205.64 ± 434.12 and -75.1 ± 364.49 respectively). Even more, we can see that both

distributions reach negative values because of the larger expected losses, indicating higher risk involved in those treatment plans as they do not prepare the landscape to disrupt future wildfire events. Values of $\lambda \geq 0.5$ lead to more compact distributions – still with negative results – by sacrificing valuable outcomes but assuring the future protection of relevant flammable areas detected by the DPV layer, mitigating the expected losses. This is reflected in the distribution of the discounted utility when $\lambda = 1$, with an expected discounted utility of 211.05 and a standard deviation of 103.45, contrasting the results above with $\lambda \in \{0, 0.25\}$.

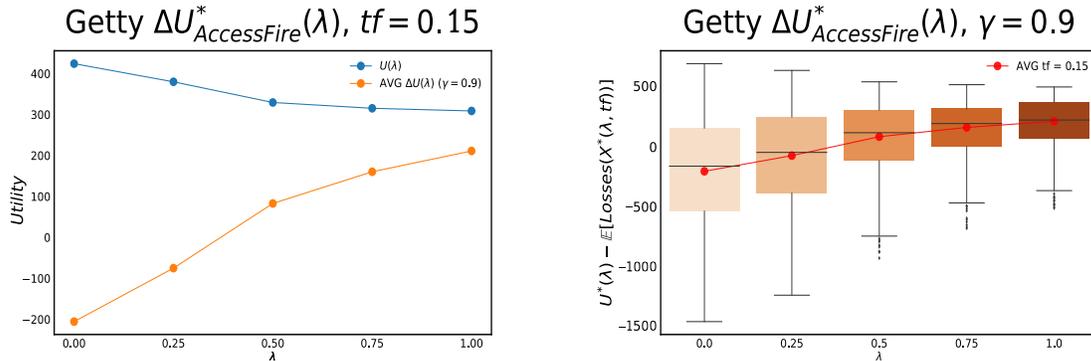


Figure 4.9: (a) Getty center instance raw utility (blue) $U(\lambda)$ with accessibility as the dominant category and average discounted utility (orange) $\Delta_{tf}U(\lambda) = U^*(\lambda) - \mathbb{E}[Losses(X^*(\lambda, tf))]$ including future expected losses due to wildfire events as a function of λ . Treatment fraction is set to 15%. (b) Distribution of the optimal utility discounted by future expected wildfire losses ($\gamma = 0.9$) for different λ levels when protecting 15% of landscape.

4.8 Conclusions

Decision-makers usually face multiple objectives when making decisions intersecting landscape and fire management. In our original approach [252], the different cells within the landscape were given equal importance. We extend these previous analyses to include multiple objectives into the optimization model and analyze the trade-off between current weighted objectives and future protection against wildfire risk. Experiments were performed on three distinct regions using utility functions obtained from the combination of demographic and environmental raster layers, highlighting different priorities of the decision-maker.

The study areas are selected to illustrate the framework when dealing with different fuel compositions and utility distributions. These differences are reflected in the Napa instance with much higher proportions of flammable fuel types having a more expanded DPV matrix compared to the other two instances presenting area-focused DPV because they are composed

of higher fractions of non-flammable landscape types. The DPV matrix provides insight into high-risk areas within the landscape and informs actions to reduce expected wildfire impacts. Regions with more compact matrices could suggest that fewer resources are needed to mitigate the effects of future wildfires in these areas. As conjectured from the calculated DPV, the expected area burned and expected losses for all utility functions as a percentage of the total instance area and the total utility available (heatmaps) was highest in the Napa instance. In addition, we observe that the fuel types commonly involved in fire propagation patterns consist mainly of some combinations of grass, woody shrubs, and shrub litter across all three instances.

In our analysis, the trade-off between present utility and the future protection of the landscape based on the DPV is evaluated. As expected, in general, the present utility tends to decrease as λ value rises as the decision-maker increasingly focuses more resources on areas with higher fire spread risk. However, present utility values are overestimated when λ drops, as the decision-maker emphasizes more the impact on the current utility function and less the future wildfire risk. Despite these common trends, certain utility layers present different patterns depending on the instance. For example, we can find that global maximum ΔU values can be obtained with $\lambda = 0$ (e.g., the forest utility layer in Napa) when the most likely propagation patterns are not associated with the highest values of the utility heatmap, reducing the risk at a very high cost of present utility. In other cases, a balanced $\lambda = 0.5$ accounts for a good trade-off between minimizing expected losses and maximizing the present utility as in the Napa U_{Equal} scenario. In addition, some scenarios including certain combinations of utility layers and instances characteristics are harder to balance with future expected losses, mainly due to the fragmentation of the instance and the distribution of the utility layer over the landscape, which can sometimes be too focused on a single point or significantly sparse, focusing the treatment plan on areas that are not likely to match the most relevant propagation patterns of the expected fires.

These results suggest that utility functions need to be carefully generated and analyzed by decision-makers in order to represent their expectations and concerns about the landscape because they play a crucial role in generating adequate treatment plans. Moreover, as mentioned in Chapter 3, the results presented in this work are conditional upon the occurrence of fires within specific time frames, meaning that the evaluation of the effectiveness of the fuel treatment plans assumes that new wildfires occur while treatments are highly effective. This may not be the case in practice, where the temporal dimension (when to apply the treatment) is crucial to implement effective plans. Some extensions to this work would involve solving a multi-stage version integrating the current framework with forest growth and the inclusion of an explicit multi-criteria optimization model including the generated utility functions as the main objectives, analyzing the trade-off of multiple feasible plans via useful techniques such as Pareto frontiers. In addition, we can incorporate wildfire severity and frequency prediction model into the framework, developing an end-to-end solution for the decision-makers. Another future direction can involve using future projected layer values (e.g., population) as model inputs to consider solutions that include projected changes in different demographic and environmental factors of interest.

Chapter 5

Wildfire suppression: A coordinated multi-agent approach

5.1 Introduction

In the previous Chapters, we have focused our data-driven models on medium/long-term problems such as landscape management and fuel treatment. In this Chapter, we exploit the characteristics and performance of the simulation model presented in Chapter 1 to explore its usage as a suitable environment to study operational wildfire decision-making problems. In particular, we focus on the operational problem where fire suppression teams are deployed after detecting an ignition and collaborative strategies are critical to contain the fire as fast as possible. Following the developments of Chapters 1 and 2, we adjust Cell2Fire to create a training environment for implementing agent-based fire suppression planning algorithms, taking advantage of the simulator speed and accuracy. However, the complexity of the wildfire phenomenon, coordination of different actors, and the interaction of several sources of uncertainty (e.g., ignition probabilities, weather forecast) make this problem a difficult challenge to solve with traditional modeling techniques (e.g., dynamic programming).

Deep reinforcement learning (DRL) algorithms have shown remarkable results [203], having solved a wide range of challenging and complex sequential decision making problems in industrial applications [105], robotics [202] or video games [344]. Extensions to natural multi-agent (MA) environments such as the coordination of autonomous vehicles [74], multi-robot control [216], network packet delivery [366] or testing environments like predator-prey, covert communication, and physical deception [212] have been developed during the last years. Despite these successful implementations, multi-agent (deep) reinforcement learning (MADRL) is an active research field: traditional algorithms and techniques developed for single-agent environments are not directly applicable, obtaining poor performance on these problems.

Although DRL algorithms have erupted in several fields and practical problems, their applications in fire management problems have not been deeply studied. Several efforts have been done to integrate fire-management decision support systems (DSS) with real/simulated

data via MIP models [356, 357, 358]. However, this approach has strong limitations: lack of flexibility and generalization (strong assumptions), intractability when adding uncertainties, and a fire-expert dependent performance. Also, the lack of well-structured links between fire simulators and planning modules remains a challenge. Despite the explosion of learning algorithms in the last years, scarce applications can be found in the fire-management literature. In [35] and [279], interactive approaches are used, updating the parameters of a system for planning the first attack once a fire is detected. However, expert dependent tuning is critical for good performance. Regarding the application of RL algorithms, a conceptual framework with applications in forest management can be found in [52], not including fire risk. The only mention of RL algorithms in the context of wildfire suppression can be found in [362] where only a theoretical framework is discussed, not providing any implementation/experiments.

In the context of multi-agents algorithms, cooperative environments have been extensively covered by different strategies and algorithms (Q-learning, Policy Gradient, Actor-Critic, see [254]) including centralized methods where multiple agents are modeled as one agent with a larger (joint) action space. Despite obtaining good results in environments with a limited number of agents, they easily become intractable with respect to their number, not being suitable for more complex environments. Therefore, centralized versions of classic DRL algorithms become non-tractable since the joint action space grows exponentially with the number of agents. This motivates the application of decentralized algorithms where agents take actions based on their local action-observation trajectories. However, coordination and credit assignment [81] problems arise in this setting leading to poor results in tasks requiring the coordinated participation of agents to reach a common global objective. Moreover, the environment becomes non-stationary for individual agents since new states are not only explained by variations of their own policy. In other words, agents are not able to quantify their unique impact on the objective function due to the perturbations of the environment, leading to poor performance.

Centralized training and decentralized execution methods where actors and critics are conditioned on local/partial observations were proposed in [147]. In this study, the contribution of individual agents to the global reward is achieved by introducing specific reward functions, triggering an effective collaboration strategy. Despite difficulties in addressing the credit assignment problem via reward functions, this work is one of the starting points for the centralized/decentralized approach. In [212], a deep deterministic MA Policy Gradient algorithm implementing centralized critics during training (one per agent) and decentralized execution of the actors is developed. Collaborative and competitive environments with continuous action spaces are tested, finding complex and efficient coordination strategies. Our approach to addressing the wildfire suppression environment is mainly inspired by [123] where a multi-agent StarCraft micromanagement environment is studied. The authors propose a centralized learning and decentralized execution approach, training a unique critic used for estimating counterfactual advantage functions in an attempt to solve the credit assignment problem [81] and estimating the individual contribution of each agent to the reward function. Using recurrent neural networks for the agents and feed-forward networks for the critics, authors obtain competitive results compared with state-of-the-art algorithms.

Therefore, we focus our research on a *centralized training and decentralized execution* approach [124, 169, 188], where agents have access to the true state of the system during training, using this information to boost their performance during the execution phase.

The main contribution of this research is the development of one of the first DSS for fire suppression planning using MADRL techniques. A novel training environment is built on the top of a state-of-the-art fire growth simulator [251], feeding agents with realistic simulated data/trajectories. The system seeks to find policies that select which areas should be harvested/treated (acting as firewalls) inside heterogeneous landscapes under fire risk in order to minimize the fire propagation damage once the fire starts. We propose a local reward extension to a state-of-the-art MADRL actor-critic algorithm [123], with the aim of finding efficient collaboration dynamics for subsets of agents with a common reward function. This approach is inspired by our environment, where subgroups of agents – usually within a certain spatial region – collaborate and split the tasks to contain the fire. Experiments in real Canadian forests and challenging collaborative instances show the potential of the system.

The rest of this Chapter is organized as follows. Section 5.2 introduces the theoretical background of single and multi-agent DRL algorithms. Section 3 presents the environment, the extensions applied to a state-of-the-art MADRL algorithm for solving our environment, the architecture used to model the agents, and describes the set of experiments. Numerical results for all experiments and relevant insights obtained are discussed in Section 4. A series of conclusions and potential future extensions are discussed in Section 5.

5.2 Theoretical Background: Deep Reinforcement Learning (DRL)

Reinforcement learning (RL) is an area of machine learning closely related to decision-making models [324]. Founded on the top of dynamic programming and Markov decision processes (MDP), reinforcement learning studies the impact of sequential actions taken over a time horizon T by intelligent agents inside an environment with the goal of maximizing an objective function, normally representing a cumulative reward over T . An environment is characterized by a set of rules/laws governing the evolution of the system through different states $s_t \in S$ as the agent interacts with it through a set of multiple actions $a_t \in A$. After an agent observes the current state of the system s_0 , it executes an action a_1 . This action impacts the environment, leading it to a new state s_1 following a transition function $\mathbb{P}(s_1, s_0, a_1)$. In its simplest form, the transitions between states follow the Markovian property (i.e., depends only on the current state and action) as in an MDP and the agent can directly observe the real state of the system before performing its action. After the transition, the agent observes a reward R_1 . This state-action-reward sequence is repeated until the end of the time horizon T , calculating the cumulative reward. RL algorithms exploit these state-action-reward trajectories in different ways to develop robust and effective

policies $\pi_\theta(s, a)$, given some parameters θ , to maximize the reward obtained by the agent in the long-term. Depending on the approach, different functions are used to approximate the expected long-term reward of the agent. The most common are the state-value $V_\pi(s)$ and the action-value $Q^{\pi_\theta}(s, a)$ functions. The state-value function $V_{\pi_\theta}(s) = \mathbb{E}[R|s, \pi_\theta]$ represents the expected reward obtained when policy π is used in state s , obtaining its optimal value $V^*(s) = \max_\pi V^\pi(s)$ across all potential policies. Alternatively, the action-value function $Q^{\pi_\theta}(s, a) = \mathbb{E}[R|s, a, \pi_\theta]$ is the expected reward obtained after performing action a in state s given a policy π_θ . Monte Carlo, temporal difference methods, Q-learning, and other methods like SARSA [324] aim to approximate these functions to find the policy (or set of policies) that optimize the long-return values, i.e., maximize the performance of the agent in the environment. Several algorithms have been developed since the formalization of RL. For a comprehensive review of RL algorithms and methods, please see [324].

DRL extends RL by using a deep neural network model to process the environment outputs (states) and mapping them to rewards and/or actions, not needing to explicitly design the state space. Thanks to this property, originally intractable problems due to the number of potential states can be addressed by reducing the state space using an approximated function provided by the neural network. One of the most known examples of DRL is the work on learning ATARI games developed by DeepMind [344], where agents trained by just passing screenshots (states) of the different games into deep neural networks can reach beyond-human performance in several games, without explicitly introducing any previous knowledge of the mechanics of the games to the agent. Given the high number of potential states (fire evolution, the position of the agents, weather conditions, etc.) and the possibility of modeling our wildfire suppression environment as a *game* where agents are rewarded when a fire is suppressed, a DRL based approach emerges as a reasonable way to address our environment.

5.3 A Wildfire Suppression Environment

We consider a multi-agent framework in a fully cooperative world, known as a Markov game [204, 254]. This is defined for a set of $n \in N$ agents interacting in an environment by performing a set of sequential actions $a_t^n \in A_n$ for each time-step $t \in T$ of the planning horizon, given a set of observations O_n and available actions A_n per agent. The true state of the environment is described by $s \in S$.

Each agent selects an action a_t^n at time-step t by following a stochastic policy $\pi_{\theta_n} : A_n \times O_n \rightarrow [0, 1]$. A joint action vector $\mathbf{a}_t = (a_t^1, \dots, a_t^n) \in \mathbf{A} = A_1 \times \dots \times A_n$ is given to the environment defining a state transition function $\mathbb{P}(s_{t+1}|s_t, \mathbf{a}_t) : S \times \mathbf{A} \times S \rightarrow [0, 1]$ describing the dynamics of the system. The initial states of the system are determined by a distribution $\rho : S \rightarrow [0, 1]$. Each agent n obtains a reward signal as a function of the state and the joint action $r_n : S \times \mathbf{A} \rightarrow \mathbb{R}$. This reward can be identical for all agents (i.e., a global reward $r_{global} : S \times \mathbf{A} \rightarrow \mathbb{R}$) or local to model different agents in the environment. Agents seek to maximize the discounted expected return $R_n = \sum_{t=0}^T \gamma^t r_n^t$ with $\gamma \in (0, 1]$ a discount factor. Global

value and action-value functions are induced by the joint action vector $\mathbf{a} \in \mathbf{A}$, calculating the expectations of R_t as $V(s_t)^\pi = \mathbb{E}_{s_{t+1:T}, \mathbf{a}_{t:T}}[R_t | s_t]$ and $Q^\pi(s_t, \mathbf{a}_t) = \mathbb{E}_{s_{t+1:T}, \mathbf{a}_{t+1:T}}[R_t | s_t, \mathbf{a}_t]$.

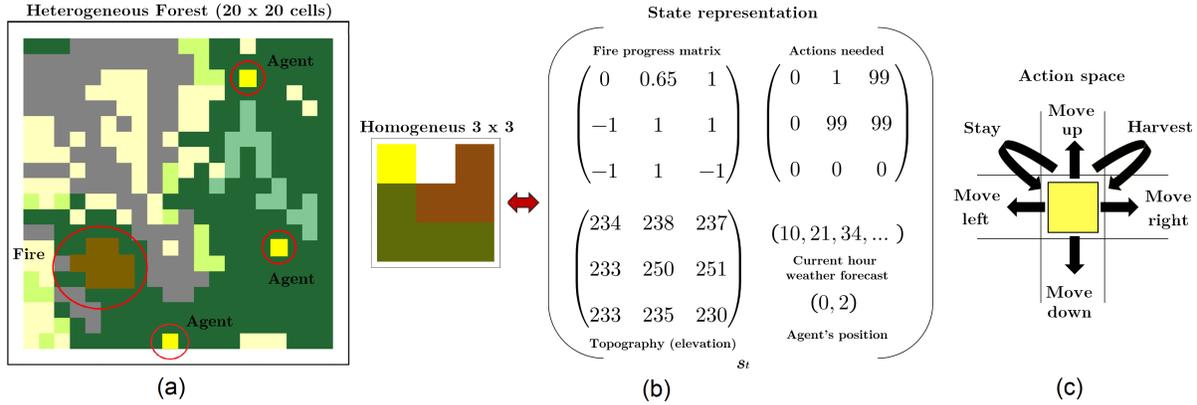


Figure 5.1: (a) Agents represented as yellow squares can harvest cells (dark green) to contain the fire growth evolution (brown). (b) The true state of the system is represented using a series of matrices and arrays potentially including the current fire progress, topographic and weather information, the position of the agents, and the number of actions required to transform a cell into a fuel break. (c) Agents can move horizontally/vertically or perform an action (stay, harvest the cell to obtain a fuel break).

Following the network formulation presented in the previous Chapters, a landscape is mapped into a grid composed of cells with an identical area where the size depends on the desired spatial resolution and data available. Each cell represents a homogeneous portion of the landscape and has two information layers: topographic and land cover. These layers define each unit, modeling them as individual objects that can interact in the landscape. We simulate fire growth with Cell2Fire by tracking the state of all cells as the fire progresses through discrete time steps $t \in T$, allowing us to retrieve that information from the environment. Using the structure of the open-source package OpenAI Gym [66], we develop a realistic fire suppression environment. The main objective of the environment is: given a detected ignition point within the region of interest, the decision maker needs to deploy N fire suppression teams with – potentially – different abilities which can take actions to modify/perturb the fuel type structure of a landscape at a cellular level to contain the fire growth dynamic as fast as possible.

A state is composed of (i) the expected fire progress, (ii) the number of actions needed to harvest a cell, (iii) topographic information, (iv) the agents' positions, and (v) the weather forecast. Agents can prioritize which cells to harvest based on how critical is the fire progress in that direction, as well as avoiding being caught by the flames. Given this information, agents could plan in advance using weather forecast data, estimating the fire growth evolution between time-steps, similar to when firefighters define their on-field strategies based on

these projections (see Figure 5.1). Agents can perform six different actions: four directional movements, harvest/treat a cell to act as a fuel break, and wait/rest. In this first implementation, all agents have identical abilities. They are penalized by a proportional factor to the fire progress motivating the agents to contain it as fast as possible, the number of burned cells by the end of the episode (-1), and being caught by fire (-100). Available cells are rewarded (+2) at the end of each episode (i.e., when no more actions are available, a team is caught by the fire or $t = T$). The environment can be naturally extended modifying certain aspects of the teams like traveling/action speed or abilities, allowing to model alternative resources such as helicopters/planes, among other possibilities.

5.4 Algorithms implemented

Deep Q-Networks (DQN)

For MDPs, Q -learning refers to a set of algorithms that aims at learning the expected long-term reward of an action in a state, i.e. $Q^\pi(s_t, a_t) = \mathbb{E}_{s_{t+1:T}, a_{t+1:T}}[R_t | s_t, a_t]$, so that an action is greedily selected to maximize the long term reward $\pi_\theta(s_t) = \arg \max_{a \in A} \{Q^\pi(s_t, a)\}$. Initial versions of such algorithms rely on building Q -tables, simple data structures containing all observed values that are updated by looping over pairs of state and action. A simple Q -learning algorithm for a single agent is to update the estimate $Q^\pi(s, a)$ with the update rule provided by [355]:

$$Q(s, a) \leftarrow Q(s, a) + \gamma(r - Q(s', a)) \quad (5.1)$$

where r is the reward observed by the agent when choosing action a in state s and s' is the state observed after choosing the action. This well-understood algorithm is rather effective. However, it does not scale well to large state spaces such as images because the size of the table grows linearly with the number of states. To overcome this challenge, [344] develops a deep Q -learning algorithm for single agent MDPs. In this work, the Q -table is replaced with a deep neural network that takes as input a vector representation of the state and outputs the estimate of the expected long-term reward of each action.

DQN estimates the expected long-term reward $Q^\pi(s_t, a_t) = \mathbb{E}_{s_{t+1:T}, a_{t+1:T}}[R_t | s_t, a_t]$ of each action. Optimizing the loss function $L_i(\theta_i) = \mathbb{E}_{s, a \sim \rho(\cdot)} [(y_i - Q(s, a; \theta_i))^2]$ with $y_i = \mathbb{E}_{s' \sim P(\cdot | s, a)} [r + \max_{a \in A} \{Q(s', a; \theta_i)\}]$, and r the reward when choosing action a in state s , s' the next state, an action is ε -greedily selected to maximize the expected long term reward obtaining a policy $\pi_\theta(s_t) = \arg \max_{a \in A} \{Q^\pi(s_t, a_t)\}$ with θ the latent parameters of the neural network.

With deep Q -learning, instead of learning $|S| \times |A|$ values, only $|\theta|$ latent parameters have to be learned. The input grows with the size of the vector representation of the state and not with the number of states. This technique made Q -learning tractable for image representation of states, for which it was intractable to learn Q -tables. When considering a multi-agent system, the size of the joint action space grows exponentially with the number of agents. Learning the Q -values associated with a state and a joint action vector is referred to

as a centralized approach. Such an approach sees its complexity increases exponentially with the number of agents for learning both the Q -tables and Q -network as the size of the output layer of the Q -network is equal to the size of the joint action space. This approach reduces to the traditional Q -learning algorithms that are proven to be efficient and will yield good policies in practice. However, their scalability is limited as they become intractable when the number of agents gets too large. To reduce the complexity of this algorithm, there have been many attempts to decentralizing Q -learning. One of those attempts is known as independent Q -learning, where each agent has its own Q -network. The idea is to let each agent have its own Q -network or Q -table to learn the Q values of the best joint actions and find a policy that selects the best of them. In [327], the authors show the gap in performance between centralized and independent Q -learning techniques and explore how shared information and mean of communication significantly improve the results.

The authors in [197] propose a distributed Q -learning algorithm to learn Q -tables in the context of fully cooperative deterministic Markov games. They restrict the reward to be non negative $r : S \times A_1 \times \dots \times A_n \rightarrow \mathbb{R}^+$. Their approach is optimistic: each agent assumes that every other agent will behave optimally with respect to its own action. By allowing the estimate $q^{(j)}$ value of agent j only to increase, they show that it converges to $q^{(j)}(s, a) = \max_{\mathbf{a}} Q(s, \mathbf{a})$. In this setting, each agent policy is deterministic and updated as follows: if the estimate Q -value for this state and action does not increase, the policy remains unchanged. If the estimate Q -value $q^{(j)}(s, a)$ increased, then the policy is updated so that $\pi^{(j)}(s) = a$:

$$\pi_0^{(j)}(s) \in A \text{ arbitrarily}$$

$$\pi_t^{(j)}(s) = \begin{cases} \pi_t^{(j)}(s) & \text{if } s \neq s_t \text{ or} \\ & \max_a q_t^{(j)}(a, s) = \max_a q_{t+1}^{(j)}(a, s) \\ a_t^{(j)} & \text{otherwise} \end{cases}$$

Moreover, the authors show that this algorithm converges to $Q_t(s, \pi_t^{(1)}(s), \dots, \pi_t^{(n)}(s)) = \max_{\mathbf{a} \in A} Q_t(s, \mathbf{a})$. However, this technique could not be directly adapted to stochastic Markov games as its convergence is not guaranteed under such a setting.

Hysteretic Q-Learning (HQL)

The work of [217] extends the previous algorithm to stochastic settings by applying a different learning rate for an increase or decrease of the estimate Q value. Similar to the deep Q -learning method, their algorithm updates the Q -values with a gradient descent step, however, the loss depends on the sign of $r + \gamma \max_{a' \in A} Q(s', a') - Q(s, a)$ where r is the observed reward when action a is taken in state s , all other action being chosen with respect to the agent's current policy, and s' the state observed after taking the action [197, 217]. Though they have no theoretical results, they show impressive experimental results on a hunter-prey problem.

Policy Gradient (PG)/Actor-Critic (AC)

The main idea behind Policy gradient methods consists of adjusting the parameters of the policy θ by directly take the gradient of the objective function $J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta(\tau)}[r(\tau)]$ and taking steps in its direction $\nabla_\theta J(\theta)$. Several algorithms have been developed, mainly differing in the way they estimate the Q -function. One of the major drawbacks of PG methods is that they suffer from high variance gradient estimates, even worse in MA environments since other agents' actions are not considered when optimizing individual policies. Different ways to reduce the variance have been documented [146] such as including a state/action dependent baseline (actor-critic), and reward-to-go, exploiting the causality of the system by noting that changes in the policy cannot affect past rewards, where we use the sum of rewards as a sample estimate of the Q function:

$$\sum_{t'=t}^T r(s_{t'}, a_{t'}) \approx \sum_{t=1}^T r(s_t, a_t) \quad (5.2)$$

In an actor-critic (AC) approach, an advantage function $A(s, a)$ is calculated by subtracting a baseline to the estimated Q -function. A critic network estimates the sum of rewards-to-go and is updated by performing a regression on target values, e.g., $y_t = r(s_t, a_t) + \gamma V^\pi(s_{t+1})$ from a target network with parameters copied periodically from the critic's network. Different approximations of the true state $V^\pi(s)$ or action-value function $Q^\pi(s, a)$ lead to variations of AC methods [324].

Difference Reward

Difference rewards methods were introduced in [365] to solve the credit assignment problem [81]. This problem appears when a global reward function is optimized in a collaborative setting and the quantification and evaluation of the contribution of each agent in the main objective are not clear since no individual signals are gathered by the agents. These methods aim to deal with failing strategies where agents are not aware of the global objective, prioritizing their own benefit instead of collaborating. Agents observe a shaped reward function D_n , the difference between the joint and global rewards when replacing the agent's action a_n by a *default action* c_n . Any action a_n that improves the value of D_n also improves the global reward function.

$$D_n = r(s, \mathbf{a}) - r(s, (\mathbf{a}^{-n}, a_n)) \quad (5.3)$$

Despite its simple but powerful formulation, the method has two main limitations: (1) depending on the environment, it is not always clear how to select/define the default action and its meaning, and (2) it requires to significantly increase the amount of simulation time since a separate simulation is required for each agent's difference reward calculation.

Counterfactual Multi-Agent Policy Gradients (COMA)

A collaborative AC multi-agent algorithm based on a counterfactual advantage function [123] estimated by a centralized critic for each agent $n \in N$ including the true state s_t of the system and the previous and current joint action vectors. It compares the Q -value for the current action a_n with respect to a counterfactual baseline where a_n is marginalized and other agents' actions \mathbf{a}^{-n} are fixed.

$$A^n(s, \mathbf{a}) = Q(s, \mathbf{a}) - \sum_{a'_n} \pi^n(a'_n | \tau_n) Q(s, (\mathbf{a}^{-n}, a'_n)) \quad (5.4)$$

The importance of this advantage function lies in the fact that it is an efficient alternative to the difference reward approach when dealing with the credit assignment problem: it does not require to perform new simulators or determine the “default” action. This approach allows to calculate all the terms in one forward pass of the critic and generating an output of size $|A|$ instead of $|A|^n$ (see Figure 5.2), capturing the impact of each agent’s decision in the global reward function. To understand this, we note that in a traditional centralized method, the number of output nodes from the network would be $|A|^n$ (all agents with identical action space A for simplicity). Thus, the size of the joint action space becomes intractable with respect to the number of agents. With COMA, the actions of the other agents are used as inputs, generating the Q -value for each action of the processed agent. Therefore, the advantage function is efficiently calculated by a forward pass of the actor and critic, generating an output of size $|A|$ instead of $|A|^n$, i.e., that scales linearly in the number of agents/actions.

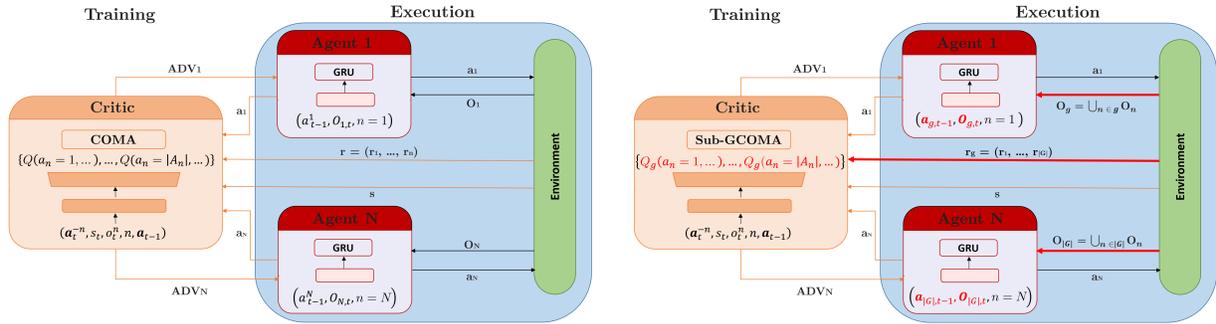


Figure 5.2: a) COMA Policy Gradient structure. A central critic computes the advantage function, used as an input for updating the actors’ policies. b) Proposed Sub-Groups COMA structure. Thick (red) lines and highlighted text indicate changes in the information passed to the agents and the critic.

5.5 A local reward extension of COMA: SubG-COMA

In this Section, we propose an extension of the discussed COMA algorithm [123] following a centralized training and decentralized execution actor-critic method with a counterfactual advantage function. We apply the concept of local rewards and local Q -functions for each sub-group of agents $g \in G$ in the context of MARL. Local rewards have been studied in the context of multi-agent reinforcement learning [36], however, they have been implemented as individual signals for each agent inside the environment and tested in simpler environments using different RL algorithms.

In our setting, each agent n belongs to a unique sub-group g , \mathbf{a}_g is the vector of actions taken by the agents in g , $r_g(s, \mathbf{a}_g)$ the local-reward obtained by sub-group g when taking actions \mathbf{a}_g on state s , and $Q_g(s, \mathbf{a}_g)$ the sub-group action-state function (Figure 5.2). Following the main structure of the original algorithm, we estimate the Q -function values for each sub-group $g \in G$ by including the groups' IDs (g) as part of the inputs of the network and keeping track of the different reward functions when estimating the *target values*. This way, a central critic calculates the new advantage function for an agent n part of the sub-group g as:

$$A^n(s, \mathbf{a}_g) = Q_g(s, \mathbf{a}_g) - \sum_{a'_n} \pi^n(a'_n | \tau_n) Q_g(s, (\mathbf{a}_g^{-n}, a'_n)) \quad (5.5)$$

Also, we allow agents from the same sub-groups to share their observations $O_n, n \in g$, increasing the collaboration degree within the same groups. Despite its implementation simplicity, this extension has a strong potential for complex environments because (1) it allows the use of different – even individual – reward function among the agents of the environment, potentially exploiting more complex collaborative behaviors while adding flexibility to the researcher when designing/testing reward functions in the environment, (2) it potentially reduces the noise when estimating the individual contribution of the agent via the advantage function, and (3) improves the collaboration performance concerning the number of agents.

Our extension is mainly inspired by the real-life operational sub-teams approach used when containing wildfires. Agents (firefighters) perform a specific collaboration with certain teams more than the others, working together in determining zones of the landscape with specific sub-objectives to reach the global goal (e.g., contain the fire in one direction to redirect it to a natural barrier or a zone being protected by other agents). This can be modeled as a set of teams obtaining a common sub-reward function that collaborates with the global reward obtained from the environment. In this case, advantage functions are calculated for the group, not the full reward function, potentially obtaining better estimations of the individual agent's contribution since we avoid calculations using the noisy global reward, both characteristics that make it difficult to assign individual credit in the original formulation.

Training Architecture

Following the structure of the original COMA algorithm [123], actors are represented by gated recurrent units (GRUs) [84] using fully connected layers for processing the inputs and outputs from the hidden state. During our experiments, we obtained similar results using feed-forward networks with three ReLu activated fully-dense layers. A similar structure is used for the centralized critic, obtaining the best results with 5 fully-connected layers of 64 units with ReLu activation. The optimal learning rate (0.0005), discount factor γ (0.99), and batch size (64) are found via derivative-free optimization methods [94, 95] using an out-of-the-box package in Python (NLOPT) similar to our approach in Chapter 2.

5.6 Experiments

Two experimental sets are used to assess the performance of the algorithms (Figure 5.3). The first set consists of real Canadian landscapes where fires on homogeneous (i.e., with identical or similar fuel-types) and heterogeneous (i.e., multiple fuel types and non-flammable cells) forests of different sizes are simulated. Homogeneous forests of 9, 400, and 1600 ha. are extracted from the Alberta province data presented in the previous Chapters, focusing on those sections with large connected areas of flammable land-cover. For the heterogeneous experiments, we use the 400 ha and 1600 ha landscape patches extracted from the Dogrib instance in Chapter 1. The second set contains five generated landscapes to assess the performance of the algorithms in specific coordination tasks. Weather conditions and ignition probabilities are based on the Dogrib publicly available forest dataset¹.

We consider several scenarios varying the number of teams $n \in \{1, \dots, 5\}$, forest characteristics, and including uncertainty sources reflected in randomized ignition probabilities and multiple weather scenarios per simulation. We test the algorithms when agents have access to the full state of the system and a partial observation of the landscape, imposing a restricted field of view (one kilometer - 10 cells radius). In the latter case, agents will only access information regarding the fire evolution and other agents in their vicinity. Policy gradient/Actor-Critic, Double Deep Q-Networks (DDQN), and Hysteretic Q-learning algorithms are implemented in their centralized and decentralized versions, comparing their performance with COMA [123] and SubG-COMA. We compare the average return of all the implemented methods, as well as the average number of burned cells by the end of an episode. Models are trained for 100,000 episodes, averaging metrics every 100 episodes.

¹http://www.firegrowthmodel.ca/prometheus/software_e.php

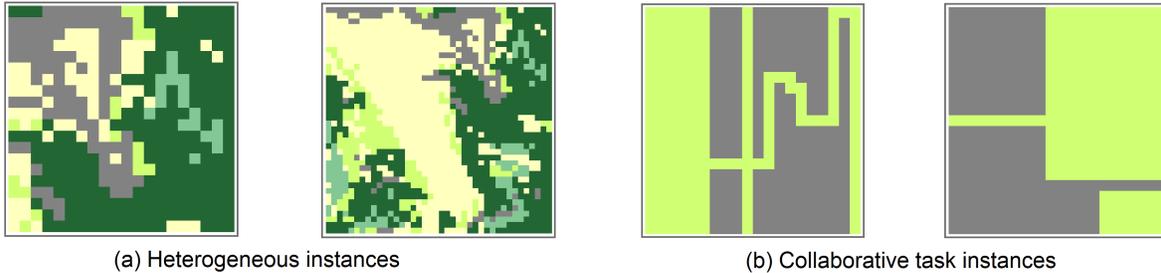


Figure 5.3: (a) Heterogeneous Canadian instances characterized by different flammable land-covers. (b) Specific collaborative tasks are tested in generated instances using non-flammable cells (grey). Given the structure of the landscape, agents require to develop complex collaboration strategies to minimize the expected area burned.

5.7 Results & Discussion

Homogeneous landscapes

Initially, we perform a series of experiments with a unique agent in the 3×3 homogeneous forest. As expected, average returns are similar across all algorithms – no coordination between agents is needed – and we obtain a trained agent that can contain the fire by trapping it inside a harvested zone. We experiment and modify the reward function, testing different penalties associated with the death of the agents. As can be seen in Figure 5.4, modifying the penalties lead to more/less cautious behaviors, giving us an important insight into the environment and learning process of our agents. In (a), all cells surrounding the ignition point are harvested to maximize the reward function while in (b) the penalty is increased, so the agent tends to avoid the section of the forest facing the main wind direction (East), minimizing the probability of being caught by the fire. Following this experiment, we test the instance increasing the number of agents up to 5 teams. All methods obtain similar performance, however, agents tend to require more time to contain the fire when no centralized training is performed, due to the lack of collaboration strategies. Moreover, as shown in Figure 5.4 (c), we observe how the variance of the returns is significantly increased concerning the number of agents when using a decentralized actor-critic algorithm – similar results are achieved with other decentralized methods – confirming our expectations.

All algorithms obtain similar performance in the 9 and 400 ha. instances, however, agents tend to require more time (+50%) to contain the fire when no centralized training is performed due to the lack of collaboration strategies. In addition, returns variance is significantly increased (+20%) with respect to the number of agents when using decentralized methods and are easily outperformed by centralized ones when dealing with larger homogeneous forests (e.g., 40×40 , Figure 5.5-(a)) where coordination is critical. Centralized methods tend to dominate in terms of average return (30% less of average area burned) but

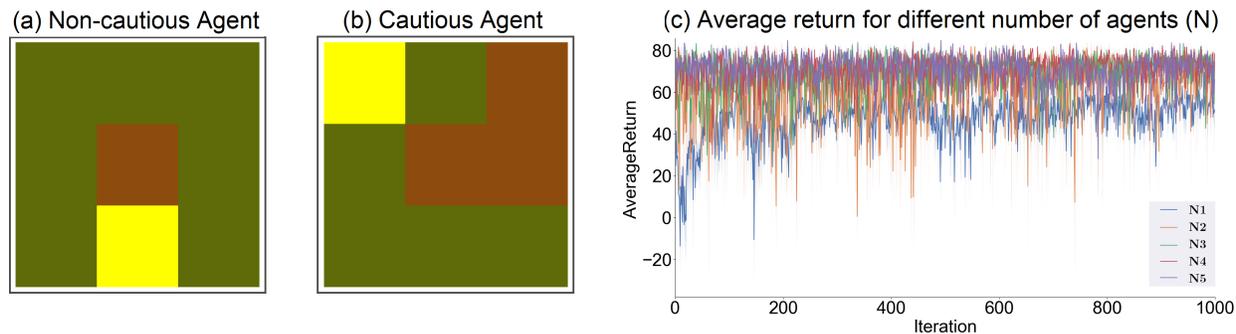


Figure 5.4: (a) A non-Cautious agent solution is able to keep the fire controlled in the ignition point. (b) The cautious agent solution avoids the eastern side of the land to minimize the risk of getting caught by the fire. (c) Average return comparison when solving the 3×3 instance with a different number of agents N , using a decentralized actor-critic method.

tend to be noisy. On average, COMA and SubG-COMA converge faster (requiring one-third of iterations) and are stable. Moreover, SubG-COMA achieves a similar performance even when we vary the number and composition of the subgroups. Q -Learning methods such as DDQN and HQL did not reach good performance and were dominated by the rest of the algorithms, thus, they are omitted from the rest of the experiments.

Similar results are obtained by varying the size of the homogeneous forest: centralized methods outperform decentralized methods but training time is significantly higher in comparison to COMA and SubG-COMA methods. Moreover, average returns are similar or even worse than the ones obtained with these algorithms. In addition, COMA and SubG-COMA present similar performance for a different number of groups $|G|$, mainly due to the fact that homogeneous instances do not require complex collaborative strategies between agents. No significant variations are observed when limiting the agents to partial observations instead of the true state of the environment.

Heterogeneous landscapes

On the heterogeneous case, we initially test the performance of a single agent in this new setting obtaining great performance with Actor-Critic based methods. In particular, COMA implementations were able to learn high-quality policies faster and with less variance than other algorithms (Figure 5.5-(a)). When increasing the number of agents, The performance of decentralized algorithms is worse than in the homogeneous case since the fire dynamic is affected by the forest structure and coordination becomes crucial to contain the fire in heterogeneous landscapes. Centralized algorithms are still competitive but they become intractable after increasing the number of agents beyond 5. Therefore, we focus our analysis on COMA and SubG-COMA methods.

Both COMA and SubG-COMA are able to learn high-quality policies faster and with less variance than other algorithms as we increase the number of agents (Figure 5.5-(b)). From the results, we observe that COMA agents tend to over-harvest the forest in comparison to SubG-COMA as the number of agents is increased. The explanation behind this pattern is that COMA agents receive a noisy approximation of their contribution to the global reward, not capturing their real impact. Thus, they perform sub-optimal actions. For example, we observe the agents trained by SubG-COMA and COMA on the 400 ha. instance (Figure 5.5-(c)) where agents 1 and 2 are located at the bottom-center of the landscape and agent 3 is placed on the north-eastern side at the beginning of the episode. SubG-COMA agents (left (C)) find an efficient collaborative strategy, using $|G| = 2$ by creating a sub-group with agents 1 and 2. This happens because the third agent observes a different reward function, allowing it to understand that harvesting cells on the northeast side is not useful to contain the fire.

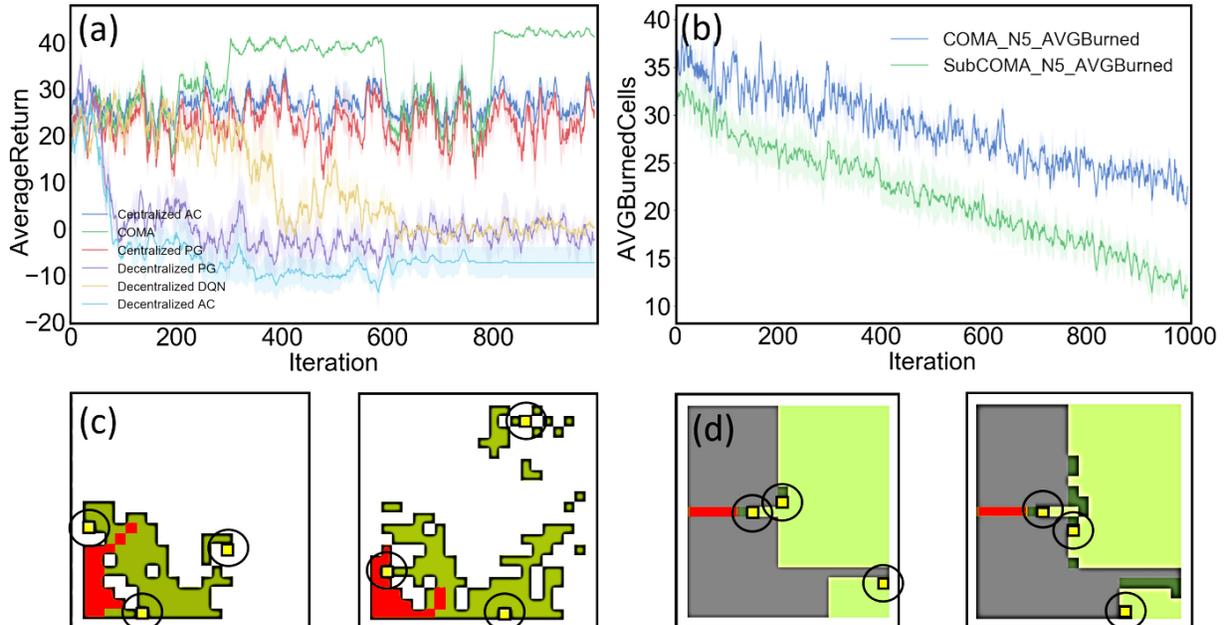


Figure 5.5: Results samples. (a) Homogeneous instance training comparison ($n = 5$, 1600 cells). (b) COMA and SubG-COMA performance comparison in a heterogeneous landscape ($n = 5$). (c,d) Visualization of SubG-COMA and COMA policies. Agents are highlighted by black circles, harvested cells in dark green, and fire in brown.

We also observe this pattern with a different number of agents trained with COMA. Increasing the number of agents leads to larger over-harvested portions of the landscape. From the experiments, we notice that SubG-COMA’s sub-groups should be carefully selected in order to exploit the characteristics/abilities of the teams. Also, understanding the likelihood

of needing a significant amount of collaboration during an episode is crucial for establishing efficient sub-grouping schemes to be exploited by the algorithm. With this in mind, extensive future research should be developed in this direction to identify if the collaboration of certain agents is critical for improving the global reward (or not) during execution instead of fixing the composition of the sub-groups before solving the environment. As in the homogeneous experiments, no significant variations are observed when limiting the agents to partial observations of the true state of the environment.

Coordination challenges

This set of instances allows us to compare and understand the potential impact of including local rewards as part of the multi-agent formulation to find effective cooperation strategies. Comparing the performance of our extension and COMA (Figures 5.5-(d) and 5.6), we see how agents following SubG-COMA are able to find subtle but more efficient/complex collaborative strategies. On the left side plots, SubG-COMA agents discover that harvesting next to the fire is enough to protect the land beyond the non-flammable section (gray cells) while COMA agents continue to harvest cells in non-risky places. We clearly observe this on the isolated agent located in the instance including an inaccessible flammable area (island instance) where the optimal action is to *wait* since it cannot help to contain the fire, however, the agent tends to harvest cells due to a failed credit assignment when training the agents while a more effective policy is achieved by using $|G| = 2$ SubG-COMA, coupling the two agents located in the mainland.

Similar results are obtained in the rest of the coordination challenges (Figure 5.7), highlighting the more effective strategies collaboration strategies obtained with SubG-COMA.

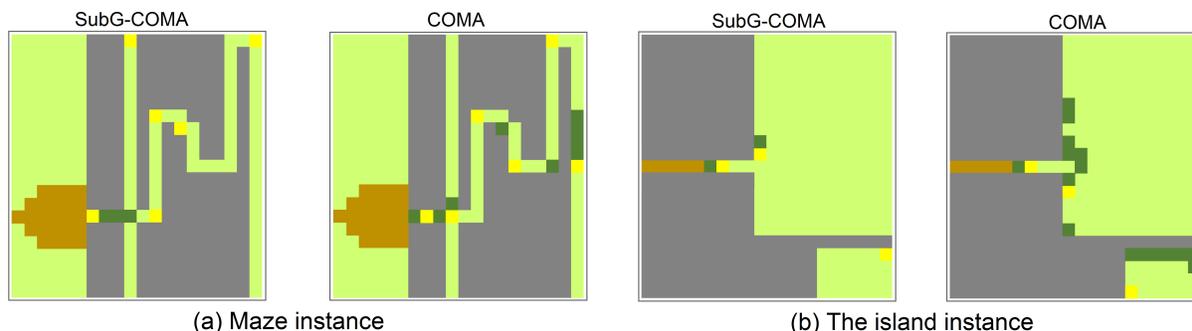


Figure 5.6: Differences between the SubG-COMA and COMA implementations when coordinating agents. Thanks to the local reward function, SubG-COMA agents are able to extract information regarding their contribution to the global reward with less noise, exploiting actions that are useful for the main objective. COMA, on the other side, is able to reach high-quality solutions but agents tend to over-harvest the landscape.

5.8 Conclusions

In this Chapter, we explored operational data-driven decision-making problems in the context of wildfire management. We focused our attention on the real-time wildfire suppression problem, where a series of agents (e.g., firefighters) are deployed to the field to control and stop the fire(s) once an ignition or fire source has been detected. Contrary to the previous Chapters, the explicit sequential and temporal dimension of the problem is translated into a complex and difficult to solve model, being intractable for traditional optimization and mathematical programming models. Moreover, the coordination of multiple agents becomes fundamental to obtain effective suppression policies.

To address this challenging problem, we explore RL and DRL algorithms, suitable for complex large-scale sequential decision-making problems. After adapting our simulation model to seamlessly interact with periodic decision-making steps, we tested state-of-the-art MADRL algorithms in a novel fire suppression environment. An extension of a centralized training and decentralized execution AC algorithm (COMA) with local rewards and Q -functions for sub-groups was implemented (SubG-COMA). We modify the reward and advantage functions of the original algorithm to provide each agent with critical information about the other teams/agents in the field, outperforming traditional algorithms in a cooperative setting.

Agents trained with this algorithm were able to find more complex (and efficient) collaborative strategies, minimizing the expected damage due to wildfires in a series of experimental landscapes. From the results, we observe that to obtain the best performance of the proposed SubG-COMA extension, sub-groups should be carefully selected to exploit complex interactions within teams, matching agents with significant collaboration.

Future work could be focused on a natural extension to multiple agents with different characteristics such as the effect of the actions/movement speed to simulate an even more realistic system, representing real situations faced by fire managers. Deeper experimentation of the proposed method could be performed to decrease the high variance as well as improving the generalization of the trained agents to multiple forest classes, via meta-learning techniques. Additionally, experiments using the proposed algorithm in other traditional (e.g. predator-prey) and challenging multi-agent environments such as the StarCraft micromanagement game would be performed in order to compare its performance with state-of-the-art algorithms in alternative environments.

In the context of wildfire management, our results represent a novel DRL application on fire suppression planning with the potential of multiple extensions and real-life applications.

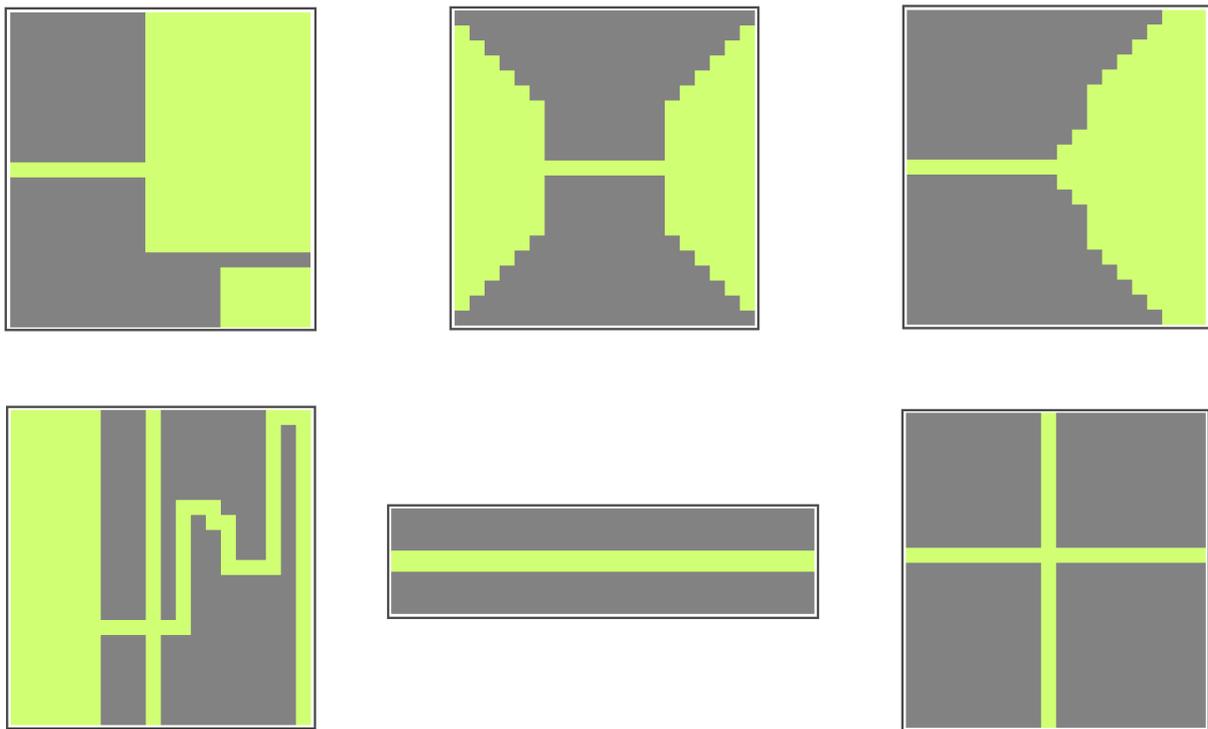


Figure 5.7: Coordination challenge instances set. Artificially generated instances are used to test the effectiveness of traditional multi-agent DRL algorithms and the proposed SubG-COMA algorithm when the collaboration of the agents plays a crucial role to effectively solve the environment, i.e., minimize the expected losses due to wildfire.

Chapter 6

Landscape planning and wildfire risk: opening the deep learning black box.

6.1 Introduction

In this Chapter, we focus on analyzing and quantifying the influence of landscape topology, understood as the spatial structure and interaction of multiple land-covers in an area of study, on fire ignition. We propose a comprehensive deep learning framework to estimate and predict the wildfire ignition risk of different areas. We focus on understanding the impact of these topological attributes and the rationale behind these results to provide interpretable knowledge for territorial planning considering wildfire ignition uncertainty. This study reveals the strong potential of landscape topology in wildfire occurrence prediction and its potential implications for the development of robust landscape management plans and similar applications using deep learning models.

Accounting for a warmer future, landscape planning could be one of the most effective ways for local adaptation to increased fire activity [61, 233]. The wildfire regime depends on and interacts with the spatial and temporal patterns of the landscapes [221]. Landscape attributes quantified by the land-cover structure, composition, heterogeneity, spatial distribution, and interactions between its components could influence fire activity, potentially perturbing multiple processes such as fire ignition, frequency, spread, the energy released, and severity in fire-prone ecosystems [132, 231, 232, 234, 346]. Given climate change scenarios, human coexistence with wildfires will necessitate informed decision-making to minimize fire risk [61, 72, 174, 233, 237]. Moreover, historical fire occurrence is an empirical manifestation of the local spatial interaction between fire and fire-prone ecosystems. Therefore, understanding these phenomena could help produce an informed fire risk evaluation and evidence-based effective territorial planning [31, 132, 226].

The analysis of spatial patterns at large scales has evolved with geographic information systems, remote sensing, and computational science capabilities [148]. This co-evolution allows multidisciplinary thinking to integrate into one framework a large amount of data from

different sources and high computational consuming analysis to solve old problems/hypotheses with new techniques [98]. In this line, deep learning (DL) models have started a revolution in multiple areas, with a significant contribution in intensive computer vision tasks such as large-scale geographic data processing and inference. DL methods are a part of machine learning techniques, a subset of the methods used in artificial intelligence (AI). DL provides a set of techniques whose mechanism is inspired by the structure of neural networks in the human brain [199] and is mathematically modeled with multiple layers of neurons. Each layer builds new variables from the previous features, which are especially useful when analyzing unstructured data, such as audio or image signals. The ability of DL to develop models at different levels of abstraction has allowed the construction of complex systems such as automatic voice recognition or object identification in images, with a performance similar to that achieved by animals and humans [199]. Interest in DL methods has grown recently owing to their remarkable success in various domains. Specifically, this interest is because of technical aspects and their effective application to problems that have not been efficiently addressed. The availability of large volumes of training data, affordable computing power, and the development of new algorithms to automatically adjust crucial parameters of the models [48] has led to the rebirth of these techniques since their early emergence in the 1980s. Their successful application in image recognition [371], speech recognition [199], text analytics [172], and self-driving cars [160], among several others, has increased the popularity of these systems.

Convolutional neural network (CNN) models have become the state-of-the-art DL models in computer vision for image analysis since their introduction in [200] and are used by researchers in many fields with remarkable success [206]. CNN-based models exploit and extract implicit information from complex spatial patterns, given their contextual-based exploration, by applying a series of filters within spatial windows. A CNN can discover crucial features by capturing multiple levels of representations from spatial data that cannot be easily detected by alternative models. Moreover, CNN models are insensitive to changes in orientation, making them particularly robust for analyzing and extracting topological features from raster data [369], understood as the spatial structure and interaction between several components of an image (e.g., the position and proportion of different pixels).

DL can be a crucial tool for numerous ecological applications owing to the complexity of ecological data and the ever-growing size of ecological datasets [86]. Recently, DL techniques have been applied in the field of ecology. Specifically, species identification [244], animal behavior classification, and biodiversity estimation for large datasets such as camera trap images [288, 289], audio recordings, and videos [311] are some examples. In the wildfire applications context, DL models trained for fire detection [316] and wildfire spread prediction [161] are some examples found in the literature [165]. However, to the best of our knowledge, only a few studies and applications of DL have been reported for understanding the effect of landscape patterns on fire activity [148], being the closest application the one presented in [368] for fire susceptibility mapping. Landscape metrics, such as proximity, connectivity, adjacency, or composition index among others, that are typically used in landscape ecology as a measure of structural characterization do not change significantly over time [148]. These

metrics are widely used with the assumption that spatial patterns can explain ecological processes and functions in a landscape. However, summarizing the spatial patterns with one or a few metrics can leave out much information, limiting the understanding of landscape patterns, interactions, and processes.

The study of fire occurrence focuses on measuring the chance that a fire might start in a region given a set of conditions/factors. Previous studies have considered risk factors derived from weather, danger indices, physiography, land-cover, or socioeconomic variables [104, 221], which are usually represented by a collection of individual (and potentially, correlated) numerical or categorical variables [132]. The machine learning field has provided useful tools and methods for modeling fire ignition. Techniques such as classification and regression trees [22], fuzzy-metaheuristic ensembles [229], support vector machines [136], random forests [215, 248] and so forth, have been successfully applied, as alternatives to traditional statistical methods. Topological aspects understood as the spatial structure and interaction between several components of the land (e.g., land-cover), are captured by constructing and extracting selected indexes of landscape structure and composition. For example, [208] includes patch density, mean patch size, mean distance to the nearest neighbor of the same category, edge density, and the Shannon diversity index for modeling the burned area in the Eastern Iberian Peninsula. Similarly, the work by [282] studies the impact of the interaction between roads and land-cover on fire ignition by constructing a selectivity index. There is sufficient evidence that the spatial configuration of these elements can influence fire regimes, especially under no extreme weather conditions [12, 221, 232]. From our point of view, this aspect of the landscapes has been challenging to study and difficult to understand due to i) the difficulty of explicitly extracting metrics that can capture aspects of the landscape topology, generally requiring more computational resources and time; and ii) the non-spatial nature of traditional statistical and machine learning models (e.g., random forest) that rely on the inclusion of multiple variables, such as those previously mentioned, to explain the phenomenon and improve the predictive performance of the model.

Land-cover data represent the basic information for landscape analysis at a pixel level of the primary land uses in a geographically delimited territory. This two-dimensional representation of the landscape also captures an intrinsic relationship between adjacency and connectivity of its components, which can be studied to understand landscape phenomena and complex patterns and processes. In the context of wildfire, effectively capturing, extracting, and interpreting this topological information can play a fundamental role in future prevention and mitigation strategies [295]. Similarly, it could inform the optimal allocation of resources across regions and the development of effective and sustainable risk management policies. Therefore, we observe an opportunity to apply new effective methods/techniques, such as DL, to try to isolate, quantify, and understand the impact of landscape topology as a predictor of fire ignition, motivating the present study. Moreover, given the current open access to huge global datasets on cloud services and the massive computational power that allows fast training cycles of the models, the scalability and potential impact of such a model are unbounded.

In this Chapter, we introduce deep fire topology (DFT), a comprehensive framework

(available as open-source software) for landscape topology analysis. Landscape topology, understood as the spatial structure and interaction of multiple land-covers in an area of study, is extracted by training a DL model while focusing on understanding the rationale behind the results, a crucial element to the applicability of the model. We present an application on the landscape ecology of wildfire to evaluate the land-cover topology as a predictor of the fire occurrence probability. We use a CNN-based model without considering additional information such as population density or weather information, which are generally used as variables [104, 132], to isolate and assess the impact of different spatial patterns in fire ignition. DFT use a supervised learning approach where land-cover images are labeled as fire or no-fire, representing the presence/absence of an ignition. We also open the DL black box by including state-of-the-art visualization and statistical techniques aimed at understanding the results. This, with the aim of informing and supporting landscape managers to minimize fire risk or improve the effectiveness of suppression strategies by optimizing the allocation of resources in high-risk areas, and potentially, support the development of new landscape managing and protection policies.

6.2 Data mining, processing, and inputs

Data can be provided to the framework using two main approaches: (i) automatic collection and processing of online assets in Google Earth Engine (GEE) to generate a ready-to-use dataset and (ii) local data provided by the user for training/prediction. The first method is integrated into our framework via scripts and can be used as an independent module for queries.

In the first approach, users can provide a set of ignition points in a latitude/longitude format, starting from our data mining and processing module. The landscape data are automatically gathered from GEE, generating a square buffer of area $A \text{ km}^2$ ($A = 1$ by default) centered at the given points. If other layers or metadata are required, for example, temperature or road density, users can execute the module independently and include any relevant source available at the GEE to be stacked as an extra band of the sampled image. Finally, the data are automatically downloaded to a cloud account or local hard drive. Although global public datasets are already available in GEE, users would need to provide and upload their high-resolution regional datasets to take complete advantage of our framework.

Alternatively, users provide unclassified data consisting of images (RGB or grayscale) with or without metadata as extra input features of the model. All samples are automatically resized and scaled to match the dimensions required by the classification model (32 32 normalized pixels by default) and are recorded locally. The dataset is subsequently used for training a custom model or processed for prediction by pre-trained models. Parameters such as the training/validation proportion or classification labels can be modified.

Study Area

We tested our approach using south-central Chile as a case study. It includes eight administrative regions ($\sim 212,000 \text{ km}^2$), corresponding to 28% of the national territory. It covers an area from 32° - 43.8°S ($2,750 \text{ km}$), which accounts for 98.5% of the national historical fire occurrence [142]. The northern part of our study area has a semi-arid and Mediterranean-type climate, with precipitations concentrated during the austral winter, a long dry season, and peak temperatures over 30°C . The southern part ($\sim 38^\circ\text{S}$) is dominated by a temperate climatic regime with higher precipitation values toward the Andean range. In the last five decades, central and southern Chile has had a highly dynamic land-use and land-cover change with different temporal and geographic variability and patterns [227]. In central (33° - 34°S) and southern (40° - 42°S) Chile, the conversion of native forest to shrublands has been the main pattern, with 45% being changed thus far. From El Maule (35°S) to Los Ríos regions (40°S), the expansion of dense exotic forest plantations has been for pulp and paper production, creating a highly homogeneous and flammable landscape [223].

Wildfire data

We used a temporal subset of the public database of wildfires provided by the Chilean Forest Service that includes the ignition point coordinates for each fire from 0.01 ha of the burned area since 1985. This database considers all type of fires, with a 99.7% human-caused fires occurrence (CONAF¹). The temporal wildfire subset includes all fires from January 2013 to December 2015. We used this temporal window because the country's most accurate land-cover map was developed in 2014 [372]. This subset includes 19,413 ignition points with a total burned area of 245,815 ha. From these fires, 70% were less than 1 ha, 23% from 1 to 10 ha, 6% from 10 to 200 ha, and only 1% over 200 ha. Among those fires, we select fires located within a distance of less than 2 km around each city, accounting for 5% of the national territory, but concentrating more than 50% of fire ignitions. The selected fires represent less than 20% of total national burned area [226] (Figure 6.1). The selected ignition points are associated with fires within this perimeter (buffer zone) to minimize landscape changes for the study period due to burned areas by wildfires.

Land-Cover data

We used Landsat 8 (OLI) derived land-cover with a spatial resolution of 30 m from [372]. These land-cover data are in a raster grid that includes ten classes: croplands, native forest, forest plantation, grasslands, shrublands, wetlands, water bodies, impervious surfaces, barren land, and snow/ice. The native forest includes primary and secondary forests of Mediterranean and temperate types. Forest plantation is an industrial tree plantation with a commercial purpose, with exotic species of *Pinus* and *Eucalyptus*. In this study area, the impervious surface is represented mainly by urban and industrial infrastructure. The

¹www.conaf.cl/conaf/seccion-stadisticas-historicas.html

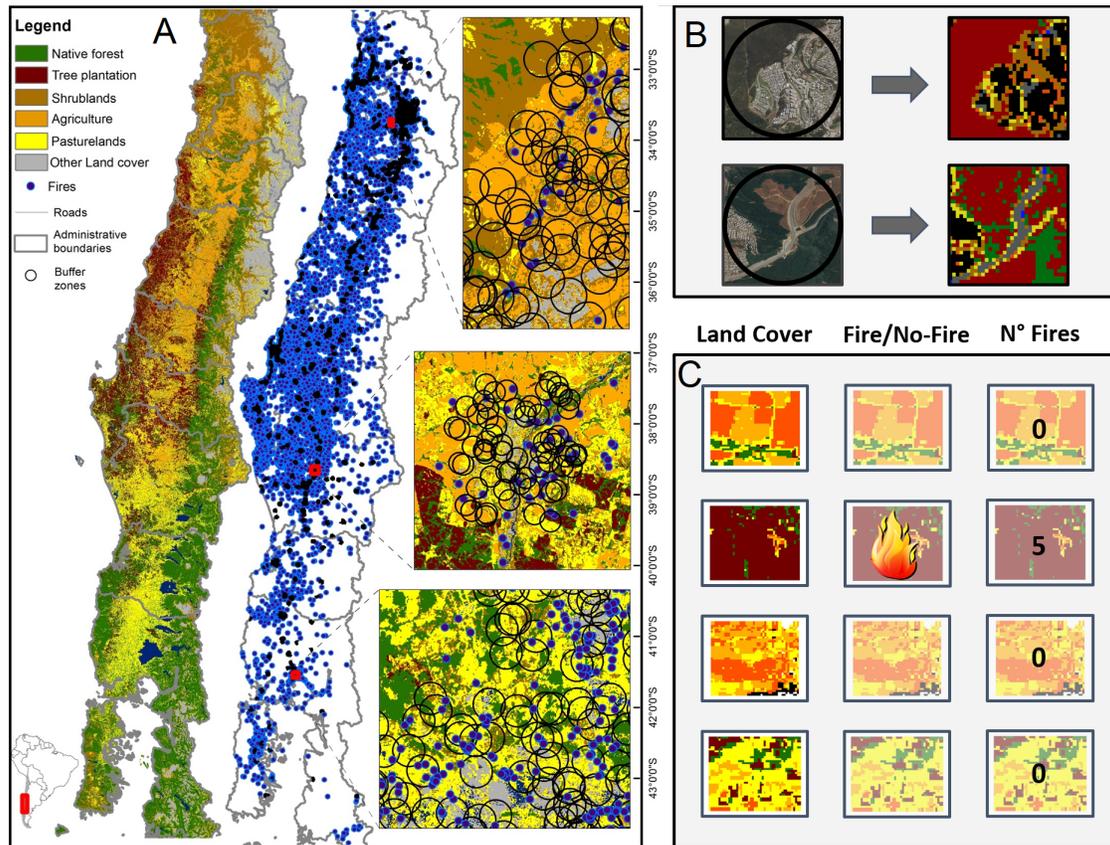


Figure 6.1: Case study data. (A) Ignition points of wildfires distributed along the national territory of Chile are highlighted. (B) A buffer is generated for each point. (C) A rectangular land-cover image centered at the ignition coordinates is extracted. The total number of wildfires experienced (numbers on the right side of the panel), their main characteristics (e.g., duration, size, and perimeters), and auxiliary variables to assess the predictions of the model (e.g., population and road density) are registered for each rectangular area as part of the dataset metadata.

landscapes are mainly composed of shrublands (29%), croplands (23%), grasslands (21%), native forests (10%), forest plantations (8%), and impervious surfaces (6%).

Wildfire and land-cover data integration

The fire point pattern association with land-cover raster data was carried out in GEE. Within a buffer of 2 km around each city, we extracted 20,000 selected square areas of 1 km² randomly centered at the generated points (Figure 6.1). These square areas could be overlapped and contain parts of the land outside the 2 km buffer around the cities. We eliminate those

square areas with a low proportion of fire-prone classes, deleting those with more than 30% of their pixels within wetlands, water bodies, barren land, and snow or ice land-covers. We associate each square area with the fire occurrence frequency within the 2013–2015 period, labeling them as fire positive (1) or negative (0) observations (Figure 6.1, (C)). Additional metadata (e.g., fire duration, size, and perimeter) are registered and associated with each area for reference and posterior evaluation/interpretation of the results, not being used as inputs of the model.

6.3 Deep learning framework

DFT is a framework and open-source software that specializes in understanding and analyzing the role of landscape topology in fire ignition by applying state-of-the-art DL techniques. Using a supervised learning approach, where land images are classified as fire positive or negative, our model predicts future wildfire ignition risk within a region of interest by exploiting the topological information of land-cover. With a significant focus on the interpretability of results, which is one of the biggest challenges in the AI field [78], the framework generates crucial outputs to understand the risk assessment process, answering the why question, and providing the logic behind the results. Together, the analysis and interpretation of these outputs allow the users to investigate (i) the propensity of spatial configurations/patterns to wildfire ignition, (ii) the creation of novel risk indexes based on topological characteristics for landscape management, and (iii) the support and development of new urban policies.

By default, classification models are trained by randomly splitting the datasets into training and validation sets, representing 75% and 25% of the images, respectively. This is naturally extended to custom models, including metadata as features, by keeping track of the image and their associated vector of numerical variables.

Convolutional neural network

We implemented a computer vision algorithm based on a CNN model for assessing the future risk of wildfire, given the land-cover raster of a region. The proposed CNN was designed and optimized to improve the prediction accuracy and performance validation metrics of the model while maintaining a reasonable size and an end-to-end training time. Like the popular Visual Geometry Group (VGGNet) architecture [314], our net stacks multiple (5) 3 × 3 filters to extract topological features on top of each other before performing max-pooling operations. Depthwise separable convolutions [85] are used instead of standard convolution layers to decrease the model complexity without significantly sacrificing its performance. These layers are more efficient because they require less memory and computation, given that the number of estimated parameters is reduced by several orders of magnitude. If used effectively, the performance is not significantly impacted, which allows for the easy training and deployment of the model with reasonable computational resources. Moreover, we include batch normalization [164] and dropout [367] layers after each block to improve convergence

and avoid over-fitting of the model. A softmax layer is used at the end of the network, normalizing the outputs to obtain P_i , the predicted probability of ignition. Higher values of P_i indicate a higher future wildfire ignition risk in the area. We performed an exhaustive grid search of the most relevant hyper-parameters of the model (learning rate, batch size, and the number of layers) to maximize its performance in our case study area.

We provide the user with a compact and easy-to-train, but powerful classification model that can be naturally extended to multiple applications.

Imbalanced data

The problem of learning from imbalanced data, where certain classes such as fire observations, are underrepresented in the dataset, has attracted the attention of both academia and industry [155]; this is also true with CNNs, where this phenomenon has a significant impact on the model performance [163].

To address this issue, our framework includes state-of-the-art sampling methods proposed in the literature (e.g., Synthetic Minority Oversampling Technique (SMOTE) [82]) and custom weights on the loss function of the model to penalize certain classification errors. In this study, we apply the SMOTE, where new synthetic observations mimicking the distribution of the original ones from the minority class are generated, increasing the proportion of these observations. In practice, synthetic examples are generated by randomly selecting an observation from the minority class, select its k (generally 5) nearest neighbors based on a distance metric (e.g., Euclidean), and create a new observation by randomly selecting an observation in the feature space of one of these neighbors and the initial observation. Moreover, we exploit the data augmentation capabilities of our DL framework. We improve the model performance by generating variations of the original images provided to the classification model. Applying a series of transformations within the context of our study, including rotation, horizontal/vertical flips, or shifts, we expand the training samples to obtain a more general and robust model.

From the original 19,413 images, we exclude those landscapes covered by more than 30% of water or unknown land-cover types. The final dataset comprises 17,579 images, with the initial ratio of positive-to-negative cases being 1:4. We use the SMOTE to account for imbalanced data in the training set (75% of the full dataset) consisting of 13,185 images with 10,521 non-fire and 2,664 fire-positive observations, respectively. Using this procedure on the training data, we obtained a balanced set of 10,567 burned (from the original 2,664 images) and 10,521 non-burned (the original negative observations) sample landscapes.

Understanding the results

A critical aspect of AI models is their interpretability [78]. Despite the development of algorithms with remarkable performance, researchers and practitioners often fail to understand the logic behind the outputs; instead, they use them as an efficient black box. This approach leads to several limitations such as erroneous conclusions, unfair/discriminative results, and

overfitted or shortsighted models focused on specific or non-relevant characteristics of the samples, among others. DL models are not an exception. Some authors have focused their attention on understanding and visualizing the outputs from different layers of the model to identify relevant patterns [328, 370].

To address this challenge in our framework, we implemented three of the most effective visualization techniques to understand the outputs from CNN networks. These techniques are: (i) gradient-weighted class activation (GradCAM) that exploits the gradients of the final layers of the network to generate attention maps highlighting the portions of the images where the network focuses for the prediction [309]; (ii) guided backpropagation [313], a technique focused on the pixel-space gradient visualizations; and (iii) guided GradCAM, combining the previous two ideas. The analysis of these outputs provides insights into the information that the model focuses on to predict the most likely class of a sample, that is, the topological configurations that are associated with fires or non-fires, the role of discontinuities within the landscape, and the proportion of different land-covers, among others.

To exploit and study the outputs generated from these visualization techniques, we implement a zonal statistic analysis procedure. Attention maps can be understood as heatmaps where darker areas represent those sections of the image with a stronger impact on the model's predictions. Conversely, lighter areas indicate that the patterns observed in those sections are not strongly affecting the outcome of the predictive model. In an extreme case, some sections of the image could not provide any relevant information (e.g., null data) to the model, completely ignoring them during the predictive phase. Therefore, we can use these heatmaps to filter the original image at different density levels, capturing and analyzing the specific topological patterns affecting the results of the model. We filter images at four α -density levels (0%, 30%, 50%, and 70%), meaning that we filter out all pixels with heatmap values below $\alpha\%$. Once an image is filtered, we compute zonal statistics including the proportion, fragmentation, and local landscape metrics within the attention maps hotspots. Thus, we can quantitatively compare the components of the landscape at different attention levels, focusing on relevant areas for the prediction while improving the interpretability of the results. Specifically, we calculated for each filtered image: the landscape composition defined as the proportion of the landscape occupied by each land-cover; the number of homogeneous patches in the area; the mean area of the individual patches; and the Simpson's diversity index to quantify the proportionality of the distribution of land occupation among the different class types within an area [220]. The analysis is implemented as a series of automatic scripts that can be used as part of our framework or as an independent custom post-processing tool, allowing the planner to define custom α thresholds to perform the analysis.

To understand and interpret the results of the model, we ordered the observations by their predicted probability of ignition (P_i), splitting them into three classes: high-risk ($P_i \leq 0.7$), medium-risk ($0.3 \leq P_i < 0.7$), and low-risk ($P_i < 0.3$). The rationale behind these thresholds is two-fold. First, they provide natural breaks in our results as major differences in spatial distribution and composition of land-cover types are observed between the images associated with those P_i values (see Section 6.5). Second, they allow us to simplify the exposition of

our results and their discussion as we compare them with previous studies using similar categories.

Evaluating the model

We evaluate the model performance by calculating conventional metrics on the validation set, including accuracy, precision, and recall, as well as by analyzing the confusion matrix, receiver operating characteristics (ROC), area under the curve (AUC), and total training time. Owing to the aim of our study, where type-II classification errors (false negatives) should be avoided, models with high recall and ROC-AUC are prioritized over high-accuracy models. To this end, we provide custom loss functions to optimize the performance of the model in specific directions.

6.4 Computational implementation

The framework is implemented in Python using Numpy, Keras, and TensorFlow libraries for the training and prediction of the CNN model. The performance metric functions are obtained from the scikit-learn package. The imblearn package provides specialized algorithms for dealing with imbalanced data. Image and raster manipulations are performed using OpenCV and rasterio packages. Cloud data gathering and processing are performed using the GEE API. Our current implementation can be found in a public repository², where we provide a step-by-step tutorial, examples, pre-trained models, and the case study dataset. All scripts can be used directly, regardless of the OS. It should be noted that each component of the framework can be used independently, thereby allowing the users to construct their dataset, train a custom model from scratch, and predict their dataset using pre-trained models.

We ran all experiments on an Intel Core i7 3.4 GHz machine with 16 GB RAM, RTX 2080 8 GB NVIDIA GPU, and Windows 10 OS. Twenty minutes of training time were required for the entire dataset (17,579 images) using optimal hyper-parameters. On average, all outputs generated at the prediction module required a total of 12 seconds per image. Pre-trained models, including weights, are significantly compact (approximately 7 megabytes) because of the proposed CNN structure.

6.5 Results and Discussion

Classification

DFT predicted the correct fire class of each landscape with an accuracy of 91.85% in the validation set after 50 training epochs (AUC = 0.98, specificity = 0.91, sensitivity = 0.96). A

²<https://github.com/cpaismz89/DeepFireTopology>

similar performance was obtained when deploying the trained model over the entire dataset (AUC = 0.98, specificity = 0.95, sensitivity = 0.95, and accuracy = 92.4%).

False-negative (186 cases) and false-positive (1,150 cases) misclassification errors only represented 5.57% and 8.9% of their classes, respectively. These results agree with our goal of achieving a high sensitivity if false-positive errors are not as harmful as false-negative cases, to prepare a mitigation strategy against the ignition risk of a landscape.

As expected, the inclusion of SMOTE and data augmentation techniques (Section 6.3) had a strong impact on the model performance, significantly boosting the original results (accuracy = 82.7%, sensitivity = 0.65, and AUC = 0.61) and improving the convergence of the algorithm from 200 to 50 epochs. A similar performance was obtained when training the model using single-channel (grayscale) or RGB images, not significantly affecting the training times or optimal hyper-parameters. Although training times were increased, no significant impact on the performance metrics was observed when training the model with larger image sizes (e.g., 64 x 64 pixels).

Analyzing the performance of the model across the three categories defined by the ignition probability P_i (Section 6.3), we observe similar performance metrics. From the results summarized in the top-right panel of Figure 6.2, we note that all performance metrics of the model tend to be stable across these groups, obtaining the best overall performance with the high-risk group (accuracy = 100%), followed by the medium-risk (accuracy = 93%), and low-risk (accuracy = 91%) categories. We observe similar values for the precision (98%, 89%, and 88%, respectively) and almost identical performance when comparing the AUC per group (97%, 100%, and 100%, respectively).

From these numbers, we note how the model can successfully learn and identify relevant spatial patterns from land-cover images to explain why a specific area, represented as an array of pixels, is more/less likely to suffer a fire ignition. With only land-cover data, DFT reaches higher accuracy than other fire ignition risk models incorporating human activity, topography, or climatic data. The authors in [226] achieved an 89.3% global accuracy using a bagged decision tree method with the same binomial response variable. In a previous study, [20] reached only a 65% accuracy fitting a logistic regression model. This ability of DFT to accurately predict risky landscapes represents a unique opportunity to understand the combination of the landscape attributes that increase the risk of fire ignition in terms of its composition, spatial structure, and interaction between different land-covers. It also reveals that fires occur in very identifiable conditions in the area of study, reflected in DFT reaching high prediction accuracy. Given this performance, the main challenge is to unravel those patterns to provide relevant insights for both practitioners and researchers.

Opening the black box

In this section, we focus our efforts on understanding the outputs of the model and the rationale behind them. For this, we try to answer some of the following questions: why an image is classified as a fire positive/negative sample, which land-covers proportions/distributions are translated into higher/lower P_i values, how are these differences being captured by some

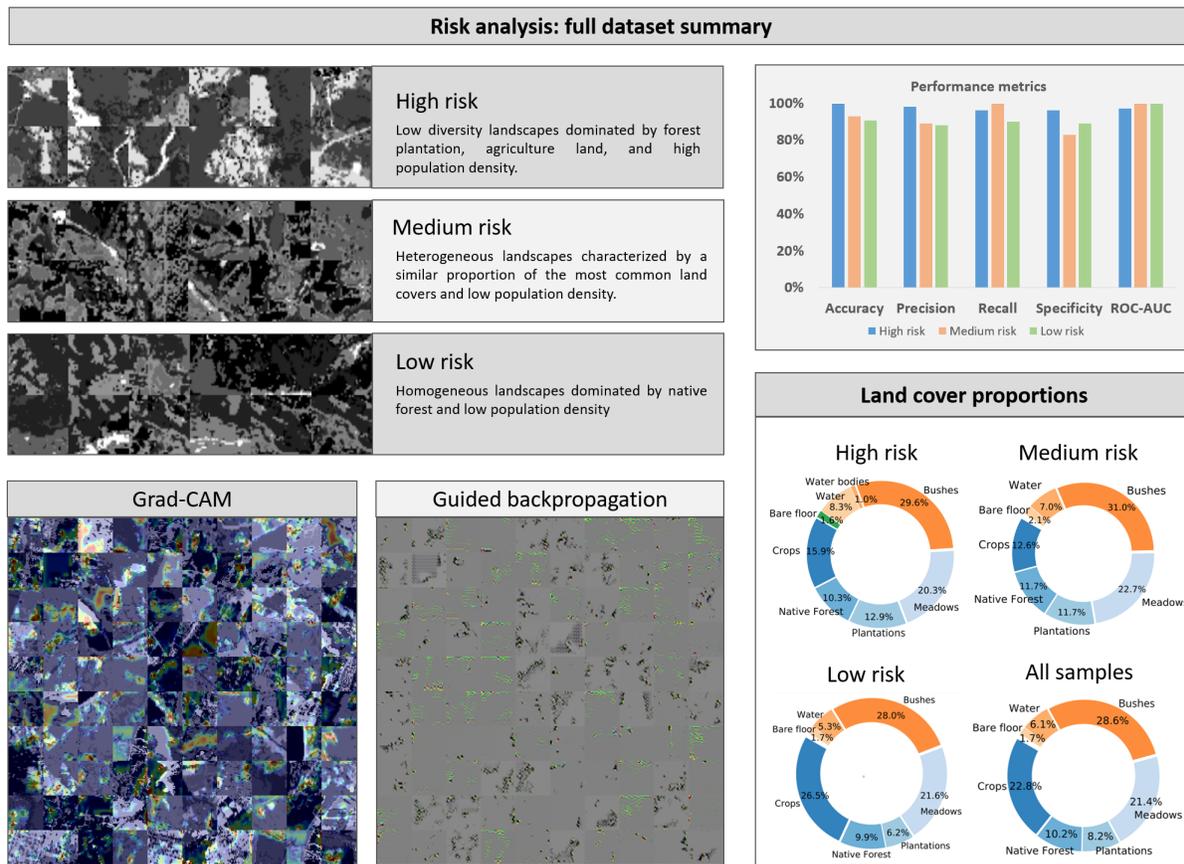


Figure 6.2: Summary results obtained for the whole dataset. Representative sample areas classified into the three ignition risk categories, defined from the predicted fire ignition probability thresholds (Section 6.3), are shown in the top-left panel. Performance metrics of the classification algorithm as a function of the predicted fire ignition probability and land-cover proportions for each group (high-, medium-, and low-risk) are shown on the right panels. Resulting attention maps from one hundred randomly selected samples using GradCAM and guided backpropagation (BP) algorithms are shown in the bottom-left panel, highlighting the areas where the model focuses its attention (warmer colors).

of the existing landscape metrics, what does the model observe when classifying an image, and how is the end-to-end evolution of the inputs during the inference process.

Focusing our attention on the outputs of the model, we can provide an initial assessment and identify the main characteristics of the images in each group (Figure 6.2, top-left panel). For interpretation purposes, we complement this evaluation using metadata associated with each image, including variables such as population and road densities. Based on the results, DFT classifies the fire-positive and negative observations accurately. Moreover, it shows

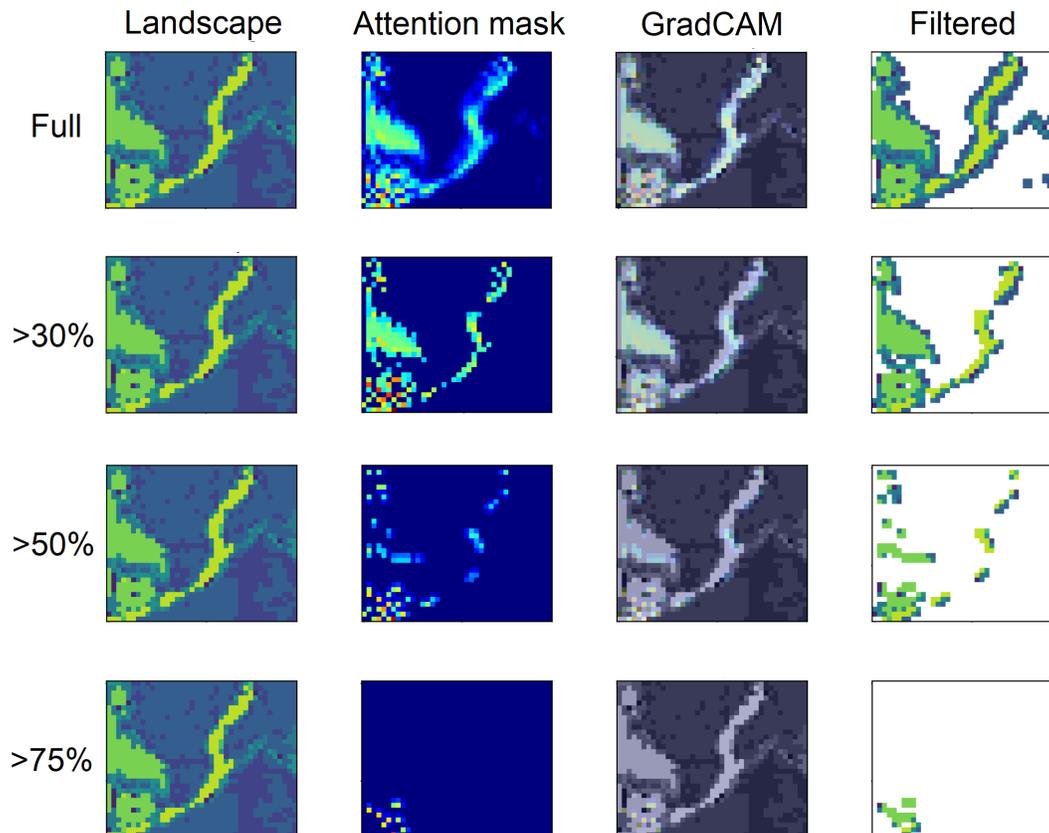


Figure 6.3: Filtered landscapes using different attention levels (full map, top 30%, 50%, and 75% of the densest areas). Attention masks are calculated at different density thresholds to filter the original landscapes, focusing the analysis on the densest sections of the attention maps. Red to blue color represent a gradient from the most to the least relevant zones to classify an image as a positive case. In this example, the network mainly highlights the urban land-cover, represented by areas covered by roads and cities, classifying the image as a fire-positive high-risk landscape.

a consistent performance when assigning the observations to the auxiliary risk categories defined as a function of P_i (Section 6.3), leading to accurate predictions within these categories. This is reflected in distinctive and characteristic proportions and spatial structures of land-covers for each group, aligned with previous studies from different regions of the world [208, 232, 318, 348] and also in Chile [20, 223, 226].

The effect of different land-cover types on fire risk in Chile has been evaluated in previous studies [141, 223, 226, 299]. However, the interaction between land-covers, continuity, and heterogeneity of the landscape concerning fire risk assessment is empirically evaluated for the first time using DL and the proposed topological interpretation framework. Observing

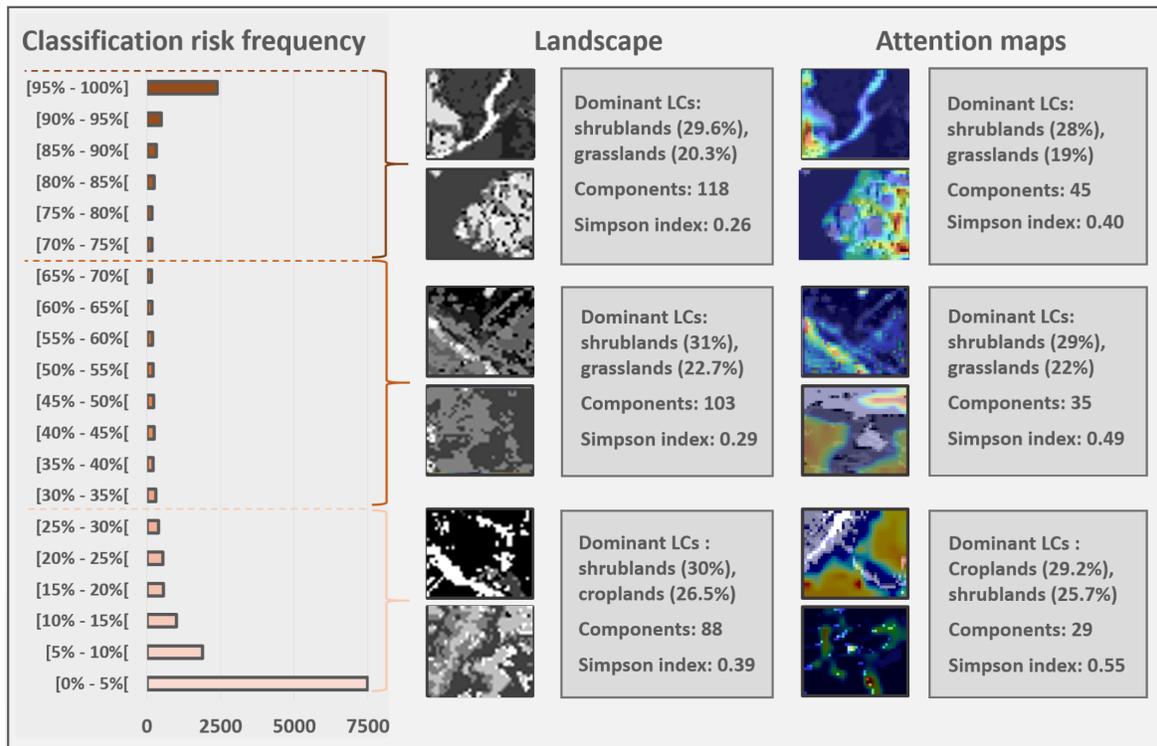


Figure 6.4: Risk predictions and characterization using activation maps. Horizontal bars indicate the number of observations at different intervals representing the probability P_i of being classified as a positive class area. Samples are then separated into three categories using P_i thresholds where significant differences are observed in terms of land-cover composition and landscape metrics. Then, these are characterized by calculating zonal statistics within the attention maps hotspots. Zonal statistics are obtained by focusing on high-density areas of the map to characterize and unveil the learning process of the classification model based on the detected spatial/topological structures and interaction between multiple land-covers. This understanding could play a crucial role in defining insightful and practical planning policies incorporating wildfire uncertainty and the mitigation of their future impact.

the results (Figures 6.2, 6.3, and 6.4), we analyzed different landscape structures and compositions to understand why each image is classified as high-, medium-, or low-risk. Images with high-risk P_i values are mainly composed of areas covered by a significant percentage of urban land-covers and forest plantations or high roads density in rural areas. They are also less diverse (i.e., fewer land-cover types per sample) than other categories, expressed in lower Simpsons index values (Table E.1). Samples belonging to the medium-risk interval present an even proportion of the most common land-covers in the study area and low population density. On the other hand, observations with low P_i values are characterized by

the presence of crops and the absence of human settlements, and proportionally, fewer forest plantations. More than 99% of wildfires in Chile are human-caused fires [142]. Therefore, fire frequency is closely related to zones with a human footprint, such as cities or other highly human-populated areas and roads, and agricultural or forest plantation industry [141, 223]. Considering this evidence, our model classified high-risk landscapes as those with land-cover dominated by industrial activities, such as forest plantation or agricultural activities, and with a high population density represented by urban areas and high road density. Conversely, low-risk landscapes are those with a homogeneous land-cover but mainly covered by native forests and represented by a low human population density. The model trained in DFT with the complete dataset coincided with these criteria. This result is significant, as it indicates that the model assigns higher or lower P_i values to those landscapes that, are known, considering the historical evidence, to be more fire-prone based on their ecological, environmental, social, or economic aspect by only capturing land-cover spatial patterns. As an example, we observe how the forest plantation cover increases from 6% (low-risk) to 13% in high-risk landscapes, in agreement with the findings in [223] and [226]; however, these landscapes have a relatively stable proportion of the other land-cover types.

Table 6.1: Zonal statistics. In this table, we present the statistics for high-risk (HR), medium-risk (MR), and low-risk (LR) images filtered with attention maps at different density levels. We include the total area of the landscape covered by the filter (% area heatmap), the average number of connected components (Ncomp), the sum of all values in the filtered landscape divided by Ncomp (MN), and the Simpson’s diversity index (Section 6.3) to illustrate our methodology.

Attention map threshold	Risk level	% Area heatmap	Ncomp	MN	Simpson
0%	HR	1	118.03	12.00	0.26
	MR	1	103.00	15.68	0.29
	LR	1	88.64	22.04	0.32
30%	HR	0.24	44.78	29.00	0.40
	MR	0.25	35.83	40.00	0.49
	LR	0.24	29.42	6.00	0.55
50%	HR	0.12	24.65	25.50	0.42
	MR	0.12	19.68	34.40	0.52
	LR	0.12	16.14	38.00	0.58
75%	HR	0.04	9.19	19.20	0.50
	MR	0.04	7.51	23.90	0.58
	LR	0.03	6.24	24.00	0.65

The detailed distribution of land-covers and their proportion within each category are

shown in Figure 6.2, bottom-right panel. Although we found some differences in land-cover composition, those differences are low for the majority of the land-covers (less than 5% on average) with the main difference being in the crops and forest plantation covers between high- and low-risk areas with 26.5% and 6.2% versus 12.6% and 12.9%, respectively. This suggests that spatial patterns or the arrangement and interaction of land-covers within the landscape could be crucial knowledge for land planning. Another differentiating pattern is the average number of patches (Figure 6.4) in the risk categories following an increasing trend from 88 (low-risk) to 118 (high-risk) components. The number of patches is not only a proxy of landscape fragmentation but also of landscape heterogeneity. This result shows that continuous and homogeneous landscape decrease the ignition risk, in contrast to, e.g., spread risk: if our response variable were the burned area or fire size, the continuity of land-cover could increase fire propagation associated with larger wildfires. Therefore, the interpretation must be cautious [148], because certain patterns may be caused by different landscape processes depending on the predominant land-cover types, climate, or anthropogenic influence in wildfire regime. Thus, practitioners should have prior knowledge of the landscape processes and dynamics for a correct interpretation of the models based on machine learning algorithms [268].

As in the analysis of the full image, insights obtained from the attention maps are consistent with the reality of the area of study, where human-caused fires represent more than 99% of fire ignitions [142]. Moreover, this reduced wildland-urban interface area (less than 2 km from cities) produces one-half of the total fires [226] mainly associated with the human footprint, expressed as human population or road density in combination with productive land-covers such as agriculture and forest plantations [140, 141]. On the other hand, the lack of these spatial patterns and the significant presence of non-flammable land-covers such as water bodies or rocky areas next to urban areas are translated into lower ignition probability values. Similarly, attention maps indicate that the model observes areas covered by large sections of homogeneous native forest patches without human presence to further decrease the estimated probability, aligned with our initial assessments. Therefore, our analysis for Chile shows that more homogeneous landscapes, reflected in fewer components and large patch areas, are associated with low-risk areas for ignitions. However, important is to note that this conclusion could be land-cover or bioclimatic zone dependent, as suggested by [223] that proposes that the preference of fire varies in a latitudinal gradient, but it is consistent with the fact that forest plantation is the preferred land-cover for fire in Chile.

For example, we can observe the attention maps generated for a high-risk landscape characterized by an urban area connected to a road and the presence of flammable grasslands and forest plantations in Figure 6.3. The classified images were filtered using the masks generated from these attention maps at different density levels. We then calculated zonal statistics (see Section 6.3) starting with the complete attention map, up to the selection of its densest area (red hotspots), characterizing the most relevant areas for the model during the classification process. From the plots, we note how both the road and city land-cover pixels are highlighted in the majority of the attention maps (warmer colors), indicating the main components of the image used by the deep learning model to estimate the ignition probability

Table 6.2: Zonal statistics obtained at different layers of the model. The table summarizes the proportions of croplands (CR), native forest (NF), forest plantations (FP), grasslands (GR), shrublands (SL), and non-vegetated (NV) covers obtained from ten random samples of each three risk categories, filtered using the full attention maps and compared with the proportions of the original landscape. Significant variations across the layers illustrate how the network focuses its attention on specific areas of the image during the classification procedure.

Attention map filtered at 0%							
Layer	Risk level	CR %	NF %	FP %	GR %	SL %	NV %
0	HR	2	9	21	17	12	38
0	MR	11	9	6	25	41	4
0	LR	0	70	17	12	1	0
8	HR	2	10	28	16	10	33
8	MR	20	15	7	20	35	1
8	LR	0	67	12	21	1	0
19	HR	1	21	43	9	7	19
19	MR	28	14	4	19	29	3
19	LR	0	57	6	35	1	1
Original	HR	1	18	47	9	6	18
Original	MR	22	13	5	22	31	3
Original	LR	0	51	13	33	2	1

of the area. As we increase the filtering threshold, we notice how the model ends up focusing its attention on the bottom-left section of the image, mainly composed of human-related land-covers and roads interacting with adjacent grasslands and forest plantations. The example also shows that the model mainly focuses its attention on the road’s insertion area in the city instead of the city or roads independently. This could suggest that increased traffic into wildlands from dense human-populated areas, as represented by this insertion point, could indicate a higher ignition risk. Similar results are observed across the entire dataset.

Despite the previously mentioned differences in general patterns between the risk categories identified by DFT, the landscape metrics that compare the complete image against the hottest attention map areas also show clear differences between the categories (Figure 6.4, Table E.1). These metrics showed that more fragmented but less diverse landscapes are classified as high-risk, which is reflected in more components and a lower Simpson’s index [220]. These differences help improve the identification of the most relevant patterns in a landscape that affect the propensity of fire ignition, having the same trend in both the complete image and the hot spots. Again, the results agree with the regions of certain human-caused fires, with fire ignition mainly associated with a highly fragmented wildland-urban interface [325].

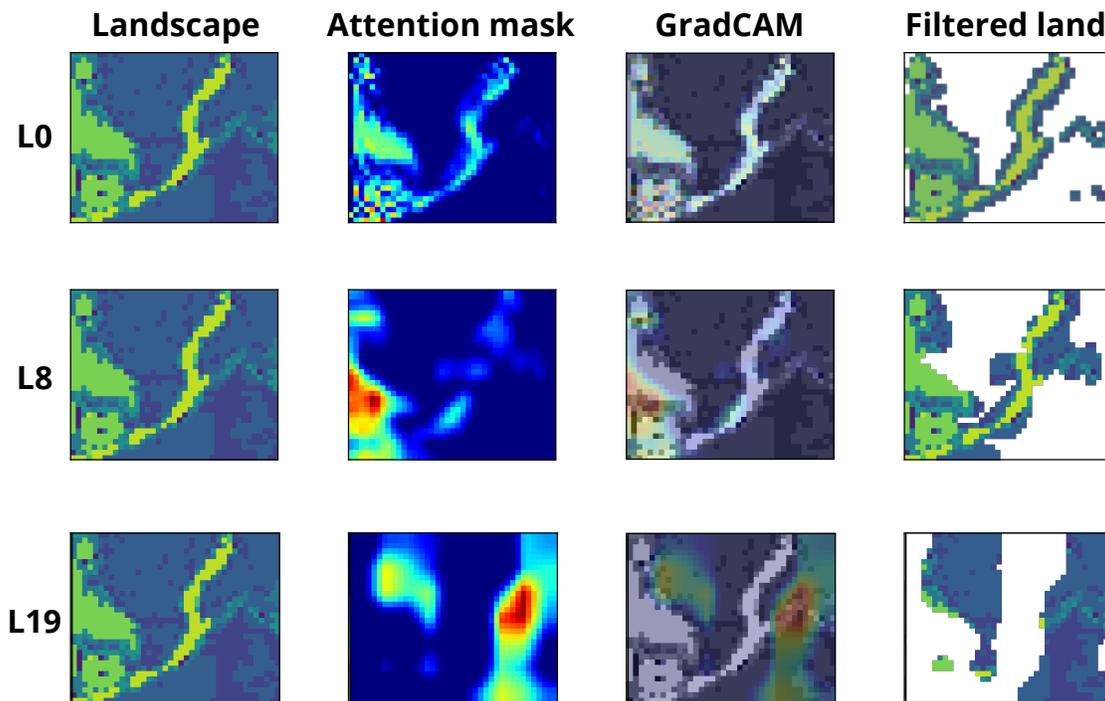


Figure 6.5: Zonal statistics at different layers of the network (0, 8, and 19). Statistics are calculated by filtering the original landscapes with the attention maps/masks generated from the GradCAM algorithm, focusing the analysis on relevant regions of the landscape. Warmer areas indicate higher levels of attention to predict the probability of ignition.

The attention maps also provide a clear assessment of which land-cover requires more attention to classify the images in the risk categories at different stages of the net, giving us insights into the learning process of the model. For this purpose, we select ten images at random close to $P_i = 1$ (high-risk), ten with P_i near 0.5 (medium-risk), and ten close to $P_i = 0$ (low-risk) to highlight the differences between landscape attributes across categories. In Table 6.2 and Figure 6.5, we show that different layers of the network (i.e., where convolutions and filters are applied) pay attention to different land-cover types. In the original image, the dominant land-cover for the high-risk (HR) image is forest plantation (47%); however, for the first layer (0), DFT focused on non-vegetated areas. The areas represented by urban land-covers and roads dominate the hottest pixels by 38% while the proportion of forest plantations is only 21%. The attention paid by DFT in non-vegetated areas even increases when filtering the landscape by the 75% hot spot, covering 73% of the hottest pixels. A middle layer of the model (8) still paid more attention to non-vegetated areas but in combination with grassland, as it is the second dominant land-cover among the hottest pixels (75% threshold). Deeper in the net (layer 19), DFT focused on the effect of forest plantations, showing patterns similar

to some recent findings [226]. The combined layers can provide a general overview of the interaction among land-covers that boosts the propensity of fire ignition. Besides, users can see where the most critical interactions occur. However, understanding the individual layer attention maps and their internal patterns opens the black box of AI, thereby unraveling different landscape patterns that could provide a deeper understanding of the ecological processes as well as support crucial management decisions in the context of landscape and urban planning policies.

To illustrate this knowledge discovery process, we follow the complete path of one image in our end-to-end classification model to understand the model functioning and help in interpreting the results (Figure 6.6). The outputs obtained at each step of the workflow show the most relevant topological attributes of the landscape for decision-making. In this case, we analyze the same HR landscape of Figures 6.3 and 6.5 composed of each risky land-cover type, such as forest plantations, human settlements, and roads. The DFT model highlights the different topological attributes of the landscape in the first convolutional layer. The results mainly highlight the highways and a combination of roads and highly populated areas represented by urban land-covers, thereby focusing not only on their presence but also their continuity because isolated pixels or small patches of roads and cities get less attention than large and continuous components. In the second convolutional layer, the model focuses on forest plantations near a town and roads, especially where roads and highly populated areas are adjacent to a forest plantation. As we move towards the final layers, the attention moves on the presence of continuous patches and their interaction with urban-related land-covers. Using this information, the model calculates P_i in its final layer, using a threshold ($P_i \geq 0.5$ by default) to determine if the sample is labeled as a fire-positive or negative observation.

Managerial insights and practical applications

The ability to live with wildfires and be adapted to their evolution is the central perspective of wildfire policy globally. For this, it is important to design landscapes that are progressively less fire-prone and more resilient. A territorial planner commanding this task generally faces two challenges depending on the state of the landscape: i) pre-fire vulnerability and ii) post-wildfire recovery. The first is generally associated with the wildland-urban interface expansion, which requires decisions on the location of homes or the construction of new roads, among others. Urban design decisions, which minimize construction costs and/or consider aspects of landscape beauty, are generally taken into account in this process without directly considering fire uncertainty, potentially producing fire-prone areas. On the recovery side, it represents a period of rebuilding physical infrastructure, realigning local institutions, and reevaluating policies that govern risk.

The decisions, in this case, are important because the first impulse (from the community and government authorities) is to recover everything as it was before the fire, thus perpetuating the problem [302]. In both cases, the new spatial configuration (e.g., of land-cover) should be evaluated. The present study and tool developed could be used for this purpose.

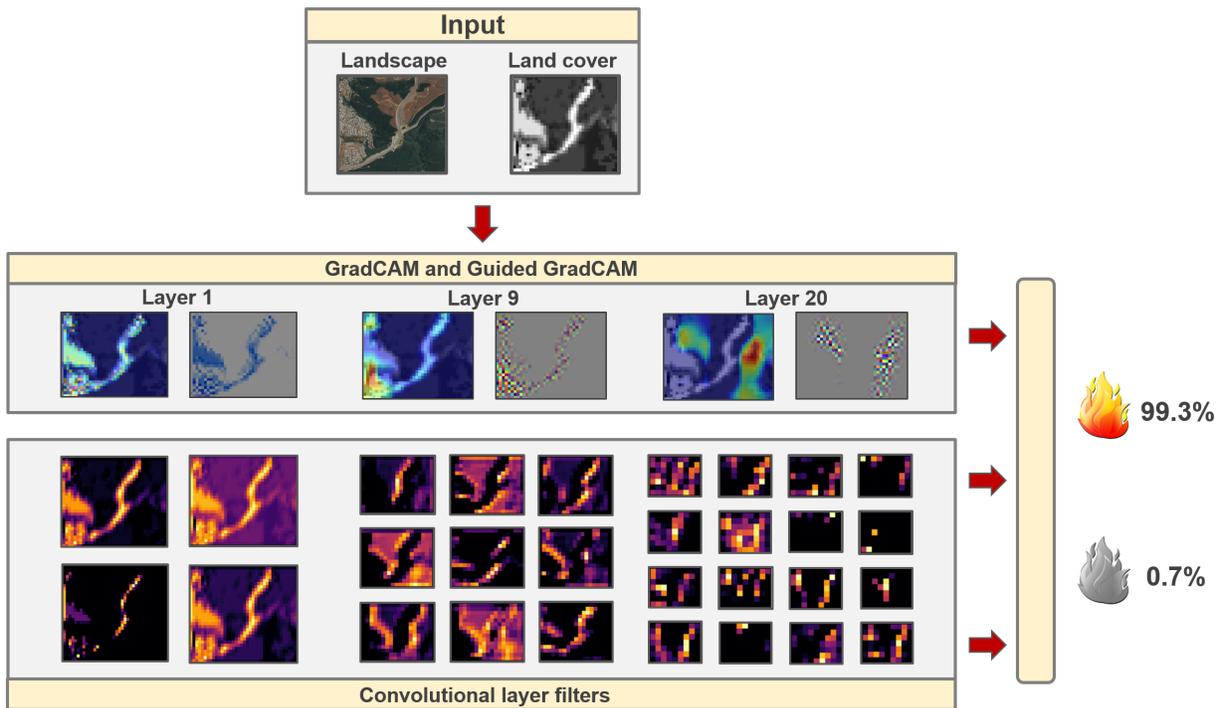


Figure 6.6: End-to-end classification of an example HR landscape. GradCAM, guided GradCAM, and filters visualizations are obtained at the different convolutional layers of the network to improve the interpretability of the model and open the DL black box. The model outputs the probability of ignition associated with each image.

The model provides an interpretable and meaningful gradient of probabilities (P_i) when processing the land-cover images, being a useful supporting tool for evaluating and guiding territorial planning activities under wildfire ignition risk. This could allow planners to identify and define custom thresholds aligned with their areas of study (with as many categories as required) and objectives, being able to exploit the outputs of our framework in similar applications requiring high interpretability of the results. Therefore, it would be possible to provide an input-map (before its commissioning) and evaluate how susceptible it is to fire ignition, identify the elements (quantitatively from the image) that are promoting the increase in risk, and complement it with other fire ignition and spread risk models incorporating relevant variables of the area under evaluation (e.g., climate) to provide a comprehensive evaluation of the proposed configuration.

On the other hand, the integration of our proposed framework with existing statistical/machine learning models to estimate ignition risk could lead to relevant practical applications. For example, we could combine the outputs of both models, weighting the generated heatmaps of the study area to evaluate different risk scenarios depending on the importance

assigned to each layer. Landscape planners could identify those areas characterized by i) the presence of risky spatial patterns, ii) dangerous combinations of variables, including e.g., climatic and sociodemographic features, or iii) the overlap of both factors. Using this information, managers could design and implement effective measures to mitigate future losses due to wildfires. These measures could include i) effective suppression resource allocation plans by identifying the most dangerous areas of the land; ii) focus the efforts on modifying those variables that can be influenced (e.g., modifying the land-cover distribution or composition in public lands) by the planner; iii) establish new policies to avoid dangerous combinations of all variables included in the models; and iv) evaluate landscape planning policy and actions in the future including trade-off of economic, social, and ecological values. Questions such as why this combination of spatial patterns, sociodemographic, and climatic variables are risky whereas the absence of certain features is translated into low ignition risk could be further study, leading to multiple research opportunities.

An open challenge

Landscape metrics measure the structural or functional characteristics of a delimited territory to quantify its spatial patterns and changes through space and time. These metrics are widely used in ecology, following the assumption that spatial patterns can explain ecological processes and functions in a landscape [337]. This pattern-process hypothesis is the basis of several landscape metrics, derived mostly from categorical maps that describe the composition and spatial distribution of each land use or land-cover [342]. However, the analysis and interpretation of landscape topology have received little attention as a novel metric of landscape patterns and as a crucial indicator to advance informed territorial planning. Also, incorporating land-cover dynamic or landscape spatial pattern change as a predictor of fire regime could expand its applications and allows a better understanding of present/future drivers and scenarios of land-cover and climate change in fire risk [231, 281, 348]. Computer vision has become essential in the analysis and extraction of relevant features and patterns from images and multidimensional data. Advances in this field have achieved remarkable results in multiple scientific studies and applications. However, interpreting and understanding the performance of models and, sometimes surprisingly and counter-intuitively learning processes, have remained an open and active research challenge [308]. In this study, we propose a flexible and customizable end-to-end DL framework integrated with automatic online data gathering and processing modules, to characterize and study the role of land topology in understanding a future wildfire ignition risk by following a supervised learning approach. Focusing on the interpretability of the results, we implement multiple visualizations and statistical techniques to understand the model outcomes.

In the field of landscape planning aimed at fire protection, this novel methodology is a step towards decreasing fire ignition risk, avoiding dangerous land-covers, their spatial patterns, and interactions. Moreover, DFT could be a valuable resource for the prevention of fires in risky areas as well as support suppression planning strategies, e.g., by allocating relevant resources according to the risk associated with the detected topological patterns of the area.

This would be reflected in more efficient and effective wildfire suppression strategies, in addition to providing supporting insights when developing landscape treatment plans (e.g., fuel treatment, prescribed fires, or thinning) that aim to minimize the impact of future wildfires. In general, however, the proposed methodology could be an essential development in landscape pattern and process interpretation in a broader sense for different ecological disciplines.

Although a step forward in the understanding of the learning mechanism of these methods and improve their outputs' interpretability, more efforts are required in this direction to continue opening the artificial intelligence black box to exploit the full potential of these powerful models.

Framework advantages and limitations

The main advantages of our framework, in addition to its performance, are its scalability, flexibility of customization, and interpretability. The seamless interaction with cutting-edge planetary-scale data cloud services allows the framework to easily scale complex and challenging global studies, enabling users to train and test state-of-the-art models without regional limitations. With thousands of high-quality datasets being available in GEE and the option of providing in-house data, the possibilities are limitless. Integrating our method with effective rebalancing methods, combined with a light CNN model that can be trained, saved, and deployed in daily-use hardware, provides our project with useful tools for various applications.

Although our example consists of a binary classification problem, the framework supports both classification (binary or multi-class) and regression models. In this regard, custom models with different research objectives can be tested (e.g., classify the landscapes into multiple predefined classes based on the type of ignition, or even by the expected area burned after an ignition) while taking advantage of the framework's pipeline. Moreover, pre-trained models provided with its code can be adjusted with new training samples by adjusting the weights of the layers.

Depending on the characteristics and similarities of the dataset and the one used during training, this method can provide high-quality results while reducing the training time. In addition, the features generated by the implemented models can be extracted as numerical variables to analyze and use them in alternative formulations, such as feeding a classification model (e.g., random forest or support vector machine), and train it with an enriched dataset in a form of "transfer learning" approach. Implementations of baseline ML models such as random forests and logistic regression, ready to exploit these inputs, are provided in the project's repository as open-source code, illustrating how our framework can interact with alternative models and methodologies. The generalization of the model to different realities (e.g., where the human factor is not as critical in fire ignitions as in our case study) is an open research question that we expect to cover in future studies, potentially incorporating extra layers of information into our framework.

Previous studies, such as [368], have effectively used a CNN model to study the fire susceptibility of China's Yunnan province by combining a variety of features (e.g., temperature, topography) with different resolutions into a unique multi-layer picture. However, this study is not oriented to understand and interpret the outputs and logic behind the proposed CNN model because of its structure, but focused on maximizing the performance of the model by including a wide range of attributes to accurately estimate wildfire susceptibility. In contrast, we focus our study on what we call learning to understand the model. We provide the possibility of observing what the model sees and understand/interpret the rationale behind the results, supported by techniques that generate relevant visualizations and statistics. Therefore, the clear pipeline, flexibility, and customization of the proposed framework allow its adaptation to multiple studies where the interpretability of results is critical to assess the performance of the model and support decision-making processes such as landscape planning under wildfire risk and the optimal allocation of suppression resources, among others. The framework is simple to use, allowing researchers and practitioners to save significant time and resources.

6.6 Conclusions

In this Chapter, we developed and tested a comprehensive DL framework in the future wildfire ignition risk assessment of landscapes exploiting topological information, defined as the spatial distribution and interaction of the data derived from land-cover maps. DFT provides state-of-the-art tools that can be used to collect multidimensional data at a large scale, train custom models, and understand the results by interpreting the model learning. It significantly decreases the training, evaluation, and model deployment cycle, and provides users with limitless and complex research questions with an efficient framework and software.

The proposed methodology adequately assesses the fire occurrence risk of a given area, using only digital information on land-cover. Furthermore, it reveals that a CNN architecture can capture drivers from the spatial configuration of landscape, showing that topology could be a key to the development of landscapes resistant to wildfire or at least, prepare them to minimize consequences due to future wildfires. This is translated into a huge potential to provide landscape managers with crucial evaluations and insights during the decision-making process. The model is able to identify and determine the risk of wildfire ignition in multiple areas by exploiting information about the interaction, continuity, and frequency of different land-covers, providing managerial insights. Given the high interpretability of the results, from both, statistical and visual perspectives, we note that the model may become a valuable resource to effectively guide fire risk mitigation and management plans, potentially informing the development of urban policies by taking into account the impact of topological patterns in the design and protection of the land.

This research can be extended in several important ways. First, the model can be generalized incorporating a series of relevant layers representing other components of the landscape (e.g., topographic characteristics such as elevation). However, this incorporation should be

performed with caution to avoid obscuring the interpretability of the results, measuring the impact of new layers in the predictions. The incorporation and quantification of these effects are relevant to a variety of applications in the fire ecology field, being an interesting future research direction.

Second, the framework could be used to generate novel landscape metrics summarizing the topological information from the image. For this, the weights of the network and their impact at different layers of the CNN could be analyzed in detail to understand the end-to-end transformation from inputs into a real-valued function describing the area of study. Third, a modified version of the framework could be used to predict other fire behavior phenomena such as the expected rate of spread, perimeter, average flame length, intensity, among other options. Depending on the application, the original framework could be complemented with analytical models, exploiting state-of-the-art formulations. Fourth, an interesting future direction could be to extend the framework to provide effective landscape management recommendations in the context of wildfire ignition risk. Given a set of high-risk landscapes, a generative model (e.g., using generative adversarial networks [101]) could be developed to provide ideal modifications in the landscape to reduce the overall risk of the area. We plan to expand our initial research in this direction in a near future.

Finally, the framework can be applied in similar computer vision studies requiring a high level of interpretability, flexibility, and customization. For this, potential modifications to the original CNN could be pertinent (e.g., number or type of layers). It could also be combined with other ML models. For example, the model could be extended to incorporate parallel deep neural networks processing different layers of information; being part of an ensemble model including other ML models such as random forests to study complex ecological phenomena; among several options. This Chapter has outlined a new DL framework in the context of wildfire ignition risk prediction, incorporating analytical and visualization techniques to understand and exploit the results of the complex DL black-box. It is, however, only the beginning of future research in this direction.

Chapter 7

Wildfires in a nutshell: a global-scale study

7.1 Introduction

Fire regimes are considered as the pivotal framing concept to understand the complex relationship between fire activity, its driving factors, and the interrelation with the ecosystems where a fire regime is allocated. Still, the different interpretations of the fire regime concept have limited the capacity to allocate specific fire regimes across the globe [191]. In order to solve such a relevant limitation, we attempt to define a methodology to frame spatially fire regimes at a global scale in this final Chapter. With this in mind, we consolidate more than 20 million wildfire records between 2000 and 2018 across the six continents. This data is processed with artificial intelligence methods to investigate and characterize global pyromes and fire regimes, areas with characteristic fire behavior over long periods of time. This spatial framing of fire regimes allows an interpretation of how a combination of vegetation, climate, and demographic features results in a specific fire regime. The current work expands on existing classification efforts by providing a unique open-source data analysis framework that makes no prior assumptions about spatial location during classification. This novel methodology bridges existing gaps between global and regional fire studies.

Fire is a global phenomenon, existing since the emergence of terrestrial plants [304] and currently present across all vegetation types [210]. The long cohabitation of vegetation and fire has induced their co-evolution [303] and shaped adaptive strategies within different plant species. Our understanding of fire activity and the relationship with its influencing factors is lacking, especially at large spatial scales [238] because of the absence of consistent long term data [236]. Currently, satellite data, despite their limited-time series length, are the primary source for analyzing fire regimes globally owing to their comparability in time and space [89]. Although studies that characterize fire activity at regional level are common [29, 92, 93, 196, 198, 235, 264, 283, 326, 347], the lack of long temporal series has limited the study and assessment of global fire activity and potential influencing factors [24, 30, 43, 49,

60, 89, 108, 109].

Several researchers have utilized global forest fire data to investigate various questions including evaluating the impact of fire on vegetation and emissions as well as studying factors influencing spatial and temporal fire activity variation [23, 152, 168, 181, 190, 238, 343]. It has been demonstrated that changing environmental conditions and human activity can and will continue to modify fire activity in several parts of the world. However, these findings have been primarily based on regional-scale studies, which do not always reflect the global reality and its variability either in influencing factors or societal perceptions [107]. The complex interrelation between environment, climate, human factors, and fire dynamics make any related study challenging on a global scale.

The relationship between vegetation and fires, for a specific ecosystem, is characterized by a fire regime [30, 51, 58, 303]. A fire regime is defined as a set of consistent and repeated wildfire conditions in a particular region over a long period of time [191, 236]. These regimes depend on a combination of factors that influence fire occurrence and behavior such as ignitions, fuel composition and arrangement, and conducive fire weather [239]. Fire conditions that define a fire regime [122, 236] have a clear impact not only on the presence of certain vegetation types but also on soil and atmospheric characteristics [89]. The combination of different factors influencing fire activity is expected to change unevenly across the globe, generating diverse patterns of the occurrence and severity of fires [181]. Although a majority of terrestrial ecosystems are familiar with fire and, in many cases, rely on its effects to maintain their natural ecosystem dynamics and health [100, 111, 276], significant changes in fire activity may lead to unwanted consequences for human communities and ecosystems conservation [60, 143, 173, 241, 276, 278]. Moreover, although the presence of a certain fire regime depends on the combination of climate, vegetation, and human activity [30], similar regimes can appear even with a different combination of factors, making future wildfire predictions a difficult task [190].

Still, the interpretation of the concept of fire regimes has changed over the years and remains a matter of discussion in the present. In [191], the authors provide insight into the discussion by dividing the present interpretations between (a) studies that consider that a fire regime or its swift reflects a core group of parameters reflecting where, when, and which type of fires occur on a certain area; (b) studies that use the factors influencing the occurrence and characteristic of fires to frame fire regimes; and (c) those that use immediate impact of fires on ecosystems and humans, as the consequence and spatial context of a fire regime. These approaches are often not used independently, and combinations between them are frequently used to provide a more comprehensive understanding of the complex relationship between fire activity, driving factors, and its impact on ecosystems [284]. However, when addressing the issue of studying fire regimes at global scales, there is a clear barrier that has impeded yet to set a specific spatial frame to those regimes, as they should combine similar characteristics in fire activity, driving factors, and impact on ecosystems. An attempt to dilute this limitation has been to define pyromes and evaluate how frequent they are within a biome [30, 89], as it is assumed that a fire regime can only be defined within a vegetation type [361]. Therefore, as science has moved from the idea that a fire regime is just a homogeneity of fire activity

over a period of time on a certain area, to accommodate new visions on influencing factors and cohabitation with a vegetation type, it has also limited its capacity to clearly delineate the extent of a fire regime.

Increased understanding of fire regimes will provide an essential tool for *knowledge transfer* between regions sharing a regime. As a result, any study on factors altering fire activity evolution or its impact will be significantly improved if implemented per fire regimes [30, 89]. In addition, the knowledge gathered could unveil useful insights and improve various studies, providing relevant information to characterize and assess the impact of current fire regimes on ecological aspects such as vegetation adaptability, soil degradation, carbon stocks, air quality/pollution, and conservation of the biosphere [50, 180, 266, 336]. Moreover, the potential evolution of the regimes, their expected impact on human communities [57, 59, 76, 175], and measures to mitigate future threats or to restore desirable conditions prior to a regime change [100] could be determined. Therefore, characterizing global fire regimes impacts multiple areas of knowledge.

Advances in generating global-scale fire data [24, 137, 138] and their accessibility [32] can be combined with artificial intelligence (AI) and processing techniques to unveil previously unseen patterns in environmental phenomena by handling large and complex data in the cloud. This is the key to better understanding and determining underlying processes in the current and future of fire dynamics across the world. To date, the use of large scale satellite data in environmental science is still restricted to satellite data [41]. Even with its current limitations, some of the premises of the use of large data [151] such as large data storage management, use of innovative processing methods to solve multidimensional problems, and sharing of data and information, facilitate the understanding of complex problems, such as identifying fire regimes across the globe.

In this final Chapter, we present a comprehensive study on investigation, understanding, and characterization of current global fire regimes using AI and a statistical framework that analyzes yearly global wildfire events over 19 years from vectorial 500m resolution datasets. We process and consolidate this data to a resolution of 1×1 degree grids covering the entire planet to calculate annual statistics on fire behavior. Next, a global classification of fire pyromes, areas with similar fire-related characteristics, is derived based on unsupervised machine learning methods. This grouping step distinguishes our analysis from existing categorization methods because we make no prior assumptions about the spatial positioning of each group. The investigation results in pyromes with significant geographical dispersion and complexity leading us to explore an extended spatial assessment of their core areas to evaluate variations in the seasonality of fire activity and to determine key underlying factors such as climate, land use, and socio-economics defining specific fire regimes.

As a result, we propose a spatial delimitation of regimes within the pyromes by defining areas with similar fire activity and select underlying factors, providing a basis for understanding the distribution of fire across the globe and the underlying factors defining them.

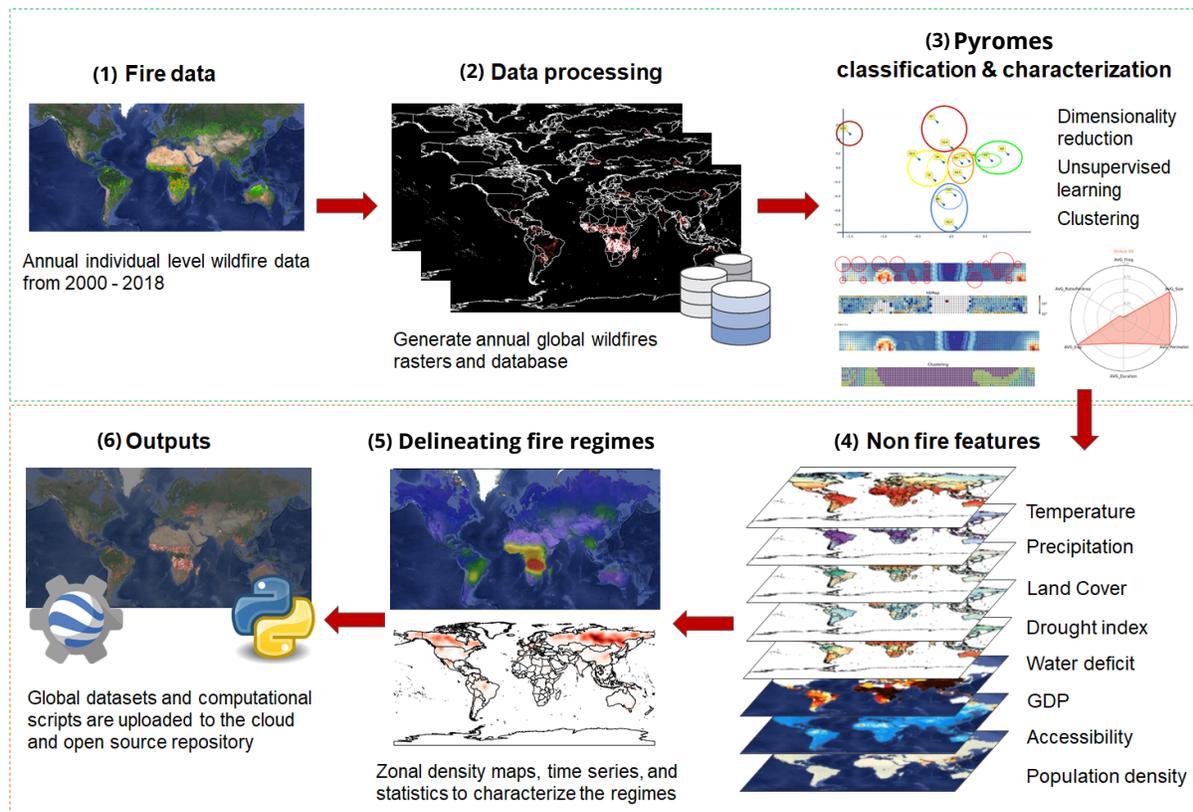


Figure 7.1: Overall framework of the study. (1) Wildfire data describing individual events in terms of fire-related characteristics such as size, perimeter, duration, and average expansion are collected from products derived from MODIS satellite observations. (2) Data are processed and consolidated into a raster dividing the world into a grid with a resolution of $1^\circ \times 1^\circ$. Annual statistics and features are calculated for each cell, generating numerical (e.g., average fire frequency per time period) and spatial datasets. (3) Statistical methods to analyze multidimensional data are combined with unsupervised learning in order to discover similar groups of cells sharing fire-related characteristics. No explicit spatial components are included. (4) Climatic and socio-economic layers are introduced for each cell in the grid. (5) Spatial density plots are generated for each pyrome, detecting the regions of the world with more observations, assumed to spacially frame a specific regime. Detected fire pyromes and regimes are characterized by climatic and demographic data. An evaluation of the influencing factors is performed for the most relevant areas. A temporal analysis to determine trends and seasonality patterns of fire activity is also carried out. (6) All results and generated datasets are deployed on cloud services and a public-access repository, along with the scripts to reproduce all steps of the study.

7.2 Global Wildfire Datasets

Two global datasets containing observations of individual wildfires in vector formats obtained from MODIS MCD64A1 collection 6 [139] with an underlying resolution of 500m were used in this investigation. This collection, although a well-known and commonly used product, is still limited. It suffers from a negative bias when estimating the total area burned at a global scale, systematically underestimating smaller burned areas, being a limitation of the dataset and our study in the MCD64A1 product. This limitation is characterized by a high proportion of commission and omission errors (40.2% and 72.6%, respectively) compared to Landsat 8 Operational Land Imager (OLI) image at a 30m resolution [55]. The lowest errors (24% and 27%, respectively) are observed in the boreal forest biome characterized by large, persistent, and well-defined wildfires while the highest error levels are achieved in the tropical forest, temperate forest, and Mediterranean biomes (> 50 and > 90%, respectively). However, both errors are found to be significantly compensated when evaluating the accuracy of the product at coarser resolutions [55] (e.g., the order of km) as the one used in our study.

The Global Fire Atlas [24] from NASA¹ provided us with about 13.3 million individual wildfires between 2003-2016. This dataset contains information about the timing and location of ignitions, the fire size, duration, and daily expansion as well as the fire line length, its average speed, and dominant direction during the spread. Individual wildfire samples between 2000-2018 were obtained from the GlobFire [32] Database². This dataset provides similar attributes to the Global Fire Atlas set (ignition, duration, size, daily burnt areas, etc.) from which the same features are obtained via simple calculations.

Although these two databases represent the best available datasets depicting individual fires and their characteristics, they have limitations. A known limitation of both datasets is the over-fragmentation of wildfire events [32], where unique wildfires are split into multiple ones due to the way these datasets are generated. In the case of Global Fire Atlas, it applies a fire event identification algorithm at a tile level (obtained from MODIS data). This approach leads to an artificial splitting of the fires following the pattern of grid tiles generated by MODIS when a fire spreads over multiple tiles. The GlobFire Database attempts to correct this limitation by introducing an algorithm that runs independently of the grid tile and source, allowing it to avoid splitting unique wildfires covering more than one tile from MODIS. Despite this improvement, it still suffers from this known limitation. Therefore, the results presented in this study are impacted by this limitation, biasing fire behavior attributes such as the average frequency and average wildfire sizes depending on the intensity of this effect in different sections of the world. In practice, we dealt with this limitation, and partially mitigate it, by consolidating both datasets. We deleted duplicated wildfires and identified as many fragmented fires as possible by comparing overlapping events between both datasets. Then, we selected the entries with the minimum number of wildfires, to minimize the impact of this limitation in the analysis.

¹https://daac.ornl.gov/cgi-bin/dsvviewer.pl?ds_id=1642

²<https://doi.pangaea.de/10.1594/PANGAEA.895835>

Finally, we generated a global raster summarizing yearly observations (2000-2018) and statistics multiple resolutions (0.05 - 1), focusing the study on the 1 product. For this, each wildfire and its attributes are projected into a grid according to the selected resolution, calculating the average frequency [fires/year], size [km²], perimeter [km], duration [days], expansion [km²/day], and the ratio between the perimeter and the area [km⁻¹] at a cellular level. A latitudinal correction is automatically applied (to all layers) to account for high-latitude area perturbations. In addition, non-flammable areas including the poles and large concentrations of water (oceans, ocean edges, lakes, etc.) are filtered from the grid to avoid bias in the analysis and comparison between regions.

Climatic, land cover, and demographic layers were obtained from global reanalysis projects providing gridded, monthly and yearly data between 2000-2018. The average, maximum, minimum 2m temperature, and accumulated precipitation at resolution of 0.25 produced by the ECMWF ERA 5 reanalysis³ were summarized to yearly data. We used the TerraClimate dataset⁴ to extract the PSDI and climate water deficit with a resolution of 2.5 arc-min. Annual land cover with a resolution of 500m was obtained from the MODIS/Terra + Aqua Land Cover Type Yearly L3 Global 500m grid (MCD12Q1) Version 6 data product⁵, that provides global land cover types yearly. It classifies the world into 17 different land covers including: water bodies, evergreen needleleaf forests, evergreen broadleaf forests, deciduous needleleaf forests, deciduous broadleaf forests, mixed forests, closed shrublands, open shrublands, woody savannas, savannas, grasslands, permanent wetlands, croplands, urban and built-up lands, cropland/natural vegetation mosaics, non-vegetated lands, and sparsely vegetated areas.

The estimated population density (number of people per square kilometer) for years 2000, 2005, 2010, 2015, and 2020 (used as a proxy for 2018) was extracted from the Gridded Population of World Version 4 (GPWv4), Revision 11 dataset⁶ at a resolution of 30 arc-s grid cell and summarized for the period under study. Accessibility to cities, measured as the land-based travel time (minutes) to the nearest densely-populated areas with 1,500 or more inhabitants per square kilometer was obtained from the Malaria Atlas Project⁷ with a resolution of a 30 s-arc for 2015. In addition, we used the gridded global datasets of Gross Domestic Product and Human Development Index between 2000-2015⁸ to extract and calculate zonal statistics for each spatial subgroup determined in our research of the average GDP per capita in a given administrative area at a 5 arc-min resolution NetCDF file, covering the period of 2000-2015. All features extracted from the different datasets were consolidated into global GeoTif files with a common resolution as well as correcting the perturbation of high-latitude regions using a Geographic information system (GIS).

³<https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels>

⁴<http://www.climatologylab.org/terraclimate.html>

⁵<https://lpdaac.usgs.gov/products/mcd12q1v006/>

⁶<https://sedac.ciesin.columbia.edu/data/collection/gpw-v4>

⁷https://malariaatlas.org/research-project/accessibility_to_cities/

⁸<https://datadryad.org/stash/dataset/doi:10.5061/dryad.dk1j0>

7.3 Clustering analysis

A numerical database derived from the global wildfire data between 2000-2018 was generated from the consolidated GeoTif files with each row associated with a cell in the global grid map. We only included cells within the land (no water/oceans). Features (columns) correspond to yearly fire-related variables. Clusters were defined using the normalized yearly average values of the whole data of the frequency (number of fires per year), size (area in square kilometers covered by the fire), perimeter (km), duration (days until the fire was suppressed), daily expansion (area evolution per day), and ratio perimeter/area of the observations, classifying each pixel of the grid as a part of a unique fire regime.

Multiple clustering and unsupervised machine learning algorithms (density-based spatial clustering of applications with noise (DBSCAN), ordering points to determine the clustering structure (OPTICS), K-Means, and self-organizing maps) along with dimensionality reduction techniques (principal component analysis (PCA) and t-distributed stochastic neighbor embedding (t-SNE)) were tested and compared. Quantitative and qualitative comparisons were performed using various performance metrics such as intra/inter distance between groups, the silhouette value [330], and the elbow method comparing the sum of squared distances from each point to its assigned center [182]. Therefore, a range of potential groups were obtained based on the described fire behavior characteristics. The final number of groups was obtained by performing statistical analysis of all possible classifications by comparing multiple subsets of features, number of clusters, algorithms, and expert assessment.

The PCA [364] algorithm is implemented to quickly visualize multidimensional data into a two-dimensional map by selecting independent components that retain the maximum percentage of the dataset variance. In our study, we implemented it as a pre-processing step before applying clustering algorithms, useful for testing the robustness of the classification, and to evaluate/visualize the incidence of the climatic and demographic components on each pyrome and fire regime. For this step, we performed an independent PCA for each set of attributes, i.e., climatic and demographic, selecting a unique component to summarize each dimension. Applying this method, we obtained unique components explaining 51% and 59% of the variance of the features, respectively. These components provide us with potential explanations regarding the driving factors behind each fire regime, characterized by the following weights: $\beta_{PDSI} = -0.135$, $\beta_{AVGTemp} = 0.694$, $\beta_{TotalPrecipitation} = 0.707$, $\beta_{GDP} = -0.590$, $\beta_{PopD} = 0.556$, and $\beta_{Access} = -0.584$.

The t-SNE algorithm [213] is a non-linear dimensionality reduction technique used for the visualization of high-dimensional datasets. It computes a probability distribution for pairs i, j of multidimensional observations such that similar objects lead to higher probability values (Eq.1, σ_i is the bandwidth of a Gaussian kernel). Then, a second probability distribution is defined over all observations in a low-dimensional map, minimizing the Kullback-Leibler divergence [194] between the two calculated distributions.

$$\mathbb{P}(j | i) = \frac{e^{(-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2\sigma_i^2)}}{\sum_{k \neq i} e^{(-\|\mathbf{x}_i - \mathbf{x}_k\|^2 / 2\sigma_i^2)}} \quad (\text{Eq.1})$$

We apply this algorithm to visualize the centroids of the regimes into a two-dimensional map, providing an effective tool for evaluating similarities/differences between clusters and interpreting the fire regimes obtained from multiple clustering algorithms. This analysis is supplemented by dendrograms providing a hierarchical structure of the centroids of the clusters (see Appendix E).

K-means [126] is used as the main unsupervised learning algorithm to classify the observations into k clusters. It seeks to minimize the within-cluster variance. We apply it to our multidimensional dataset with different pre-processing techniques (e.g., normalization, standardization, and dimensionality reduction methods). Although this is an NP-hard problem [19], it can be solved with heuristics converging to local optimal values.

K-mean models are fitted for all $k \in [2, 35]$. Performance metrics (e.g., silhouette, the sum of squared distances) and all centroids statistics are recorded and a series of plots are generated to visualize the core characteristics of each fire regime. From the experiments, convergence in the performance metrics is observed with $k \in [10, 20]$ and significant improvements are observed until $k = 10$. We determine the configurations that best balance the trade-off between complexity, interpretability, and reproducibility. Once a cell is classified, a label feature is generated in a GeoTif band, capable of remapping the cells into the world grid for visualization purposes and expert assessment. A Gaussian kernel is applied to smooth the results.

We compare the results of the previous algorithm with DBSCAN [301], a non-parametric clustering algorithm that groups points closely packed together (based on a distance metric) and can detect outliers points (belonging to low-density regions) as noise, discarding them from the grouping procedure. The algorithm finds the set of points in the ε -neighborhood of point x , i.e., $N_\varepsilon(x) := \{y \in \mathbb{R}^n : \|x - y\| \leq \varepsilon\}$, where ε is a parameter provided by the user (100 in our study). Next, we determine core points which are points with $|N_\varepsilon(x)| > \text{min}P$ where $\text{min}P$ is the minimum number of points needed to be considered as a core point (35 in our study). Using a network approach, connected components of the core points are determined. Finally, non-core points are assigned to their closest ε -cluster satisfying the cardinality constraint, otherwise, they are labeled as noise. As opposed to K-means, it does not require to specify the number of clusters k in the data before training and there is no assumption about the shape of the clusters or their distribution. Similar results are obtained using the OPTICS [27] algorithm.

Self-organizing maps [184] correspond to artificial neural network models applied in the context of non-linear dimensionality reduction of high-dimensional data. Following an unsupervised learning approach, models are trained using competitive learning, where nodes of the network are activated (compete) to represent certain characteristics of the data. Training samples are fed into the network, calculating their Euclidean distance with respect to the weight vectors of the nodes, initialized by sampling from the subspace generated by the two largest principal component eigenvectors [90]. The weights of the neuron(s) with the closest distance (best matching units) are updated using Eq.2 where $W_u(t)$ is the weight of neuron u at time-step t , $\theta(v, u, t)$ is the distance function between neuron v and u at step t , $\eta(t)$ is

the learning rate, and $D(i)$ is the input vector of the i -th training sample.

$$W_u(t+1) = W_u(t) + \eta(t)\theta(v, u, t)(D(i) - W_u(t)) \quad (\text{Eq.2})$$

This process is repeated for each input vector for a number of training epochs or until convergence within a small $\delta > 0$ threshold is achieved. Once the map is obtained, the components of the U-Matrix [341] are calculated as the average distance between the vectors of node's weights and their closest neighbors (eight nodes in a rectangular grid), representing high-dimensional data as a two-dimensional matrix. By applying a color gradient (e.g., blue to red), the U-matrix can be used to determine potential clusters of data where cooler sections (compact regions) can be considered as clusters and warmer areas represent the boundaries (larger distances) between these groups.

This representation can help to visualize the clusters in high-dimensional spaces or to recognize them automatically using relatively simple image processing techniques. We process the resulting U-matrix with a Laplacian of Gaussian (LoG) [185] kernel, one of the most widely used and effective blob detection methods aimed to detect homogeneous areas within an image that are significantly different (e.g., color and brightness) from their neighborhood. Belonging to the family of differential methods, based on derivatives of the function with respect to the position, it consists of a sequential application of a Gaussian kernel (to smooth the image and attenuate its noise) and the Laplacian filter, emphasizing regions of rapid intensity change. Thus, the LoG operator (Eq.3) used to convolve with the U-matrix is obtained to determine potential fire regime classifications.

$$LoG(x, y) = -\frac{1}{\pi\sigma^4} \left[1 - \frac{x^2 + y^2}{2\sigma^2} \right] e^{-\frac{x^2 + y^2}{2\sigma^2}} \quad (\text{Eq.3})$$

7.4 Spatial and temporal analysis

Gaussian kernels using a radius of 5 and bandwidth h that minimizes the mean integrated squared error (Eq.4) – measuring the difference between the original function $f(x)$ and its kernel density estimator $\hat{f}_h(x)$ – are applied for the spatial characterization of regimes. We determine the regions of the world with the most fire observations, based on the density of cells belonging to a particular regime. Contour lines are calculated for each local region (regime) accounting for 10, 30, 50, 70, and 90% of the local observations to determine the areas of the world where the fire regime is focused. Regions with at least 30% of the local observations are then ordered by area (largest to smallest), characterizing the top five or maximum numbers with a significant area in terms of demographic, climatic, and soil features. Once the regions are determined, historical time series are generated for each regime for all relevant features, focusing on the inter-annual and monthly aggregated fire frequency and average fire size, to determine regional seasonality and trends.

$$MISE(h) = \mathbb{E} \left[\int (\hat{f}_h(x) - f(x))^2 dx \right] \quad (\text{Eq.4})$$

The time series are decomposed into trend (T), seasonality (S), and noise (ε) using a linear additive decomposition model where $y(t) = T(t) + S(t) + \varepsilon(t)$. Seasonal components are filtered using a centered moving average. The average of this smoothed series, for each period, corresponds to the seasonal component.

7.5 Identifying Pyromes and Fire Regimes

The proposed framework is presented in Figure 7.1, emphasizing the knowledge discovery and extraction process of the study. Core operations are performed at the third stage of the framework, where unsupervised learning algorithms and techniques to process multidimensional data, such as self-organizing maps (Figure 7.2), are applied to determine patterns from the data and to validate the interpretability of the obtained results.

First, cells covering the world are classified into different pyromes based on historical fire behavior and observed patterns within their location. The observed fires are characterized by six variables: the average frequency [fires/year] (0.30 to 1175.74), size [km^2] (0.28 to 511.61), perimeter [km] (2.25 to 102.79), duration [days] (1.34 to 12.81), expansion [km^2/day] (0.23 to 18.23), and ratio between the perimeter and the area [km^{-1}] (0.78 to 8.45). We determine fifteen clusters defining general fire pyromes (Figure 7.3) distributed across the globe (Figure 7.4). These pyromes significantly differ in the mean characteristics of their fires and their spatial distribution (see Appendix E for details per pyrome and regime).

These pyromes can be further condensed into six relevant macro-groups sharing clear fire behavior characteristics (Figure 7.3): very large, fast-spreading, and frequent wildfires (R10); large and frequent fires (R11, R2, and R4); medium-sized, slow-spreading and infrequent fires (R7, and R14); small, medium-to-high frequency, and long-shaped fires (R13, R0, and R5); small and infrequent fires (R9, R1, and R8); and small/medium and very frequent wildfires (R12, R3, and R6).

Further analysis of the different pyromes based on their fire sizes, frequencies, and expansion rates reveals that R10 regime, mainly distributed across northern Australia and Southern parts of Africa, is defined by the largest and fast expansion fires, with an average fire size of 511 km^2 and $18 \text{ km}^2/\text{day}$ average expansion. Following in ranking of the fire size but with medium-to-high occurrence frequencies (average of 172 fires per year), are R11, R2, and R4 pyromes with mean fire sizes of 107, 34, and 24 km^2 , respectively. These four pyromes are often spatially associated with each other and the region covered by R10 pyrome. We observe that R11 pyrome regions usually surround the most fire-affected R10 areas. In some cases, the surrounding R11 cluster is accompanied by milder R4 and R2 pyromes, even though the latter pyromes generally occur in the Central African region, Brazil, and Kazakhstan. This spatial pattern of the observed fire activity matches with the gradient of environmental conditions, a common process of several ecological phenomena [183].

Pyromes R7 and R14 have average fire sizes of 34 and 9 km^2 , respectively, but significantly lower frequency than the previously mentioned. The occurrence of R7 and R14 show similar patterns, being mainly distributed across the boreal forests of America and Asia.

A heterogeneous macro-group was created from R13, R0, and R5 pyromes, consisting of small-to-medium fires, with sizes between 2 to 5 km², and medium frequencies. In the case of R0 and R8, there are 31 and 9 fires per year, respectively. In the case of R13, there are 307 fires per year on average. These pyromes are distributed across most of Europe, Asia, and America, on warmer zones than those where R7 and R14 are distributed.

R12, R6, and R3 pyromes comprise a more spatially compact category defined by small to medium sizes (2.94 to 4.7 km²) and very high frequencies (598.62 fires per year on average) of fires occurring almost exclusively in the tropical areas of Africa and in South Asia. Finally, regimes R8, R1, and R9, with average fire sizes of smaller than 2 km² and low frequencies (lower than 2 fires per year for R1 and R9, and 9.44 fires per year in the case of R8), are widely distributed across the world, particularly abundant in both cold and dry vegetated regions as well as in wet evergreen tropical forests. From the above mentioned analysis, it can be seen that the characteristics of fires and their distribution confirm that there is a relationship between vegetation, climate, socio-economic factors, and fire activity. However, the distribution of the areas corresponding to a majority of the fire pyromes is rather dispersed (Figure 7.4), and no single combination of factors seems to explain the occurrence of individual regimes.

Next, we deepen the analysis by exploring the spatial distribution of the pyromes across the globe in an attempt to define fire regimes. For this, we determine the most relevant (largest and densest) disjoint regions belonging to each pyrome. Fire regimes are determined by a spatial density analysis (see Section 7.4), indicating the regions of the world where observations belonging to each pyrome are mainly distributed. We denote the regimes by the combination of the regime alias with letters (a, b, c, d, and e) associated with their area (decreasing order), i.e., R1-a denotes the largest regime of pyrome R1 and, R2-c the third-largest regime belonging to R2 (see Figures 7.5 and 7.6). Interestingly, although the regimes within a pyrome share historical fire patterns, significant differences can be observed for those regimes in terms of location, climate, and socio-economic variables (Figure 7.5) and the proportion of dominant land covers. Here we include socio-economic descriptors (per capita gross domestic product (GDP) [USD], population density [individuals/km²], and accessibility [minutes]) and climatic conditions (historical temperature [C], precipitation [m], water deficit [mm], and Palmer drought severity index (PDSI) [18]) Similarly, different seasonal patterns and trends of fire-related variables such as frequency and average size, are observed from a time series analysis within the areas of different regimes as represented for the five largest regimes related to pyrome R1 in Figure 7.6.

When evaluating the relationship between the location of fire regimes within the identified pyromes and potential driving factors (i.e., climate and socio-demographic conditions), different patterns and trends can be observed. However, due to their significant variability, there is no conclusive evidence to identify the main factors defining each pyrome. Therefore, we proceed to analyze the regimes within each pyrome. By applying an independent principal component analysis (PCA, see Section 7.4) on the climatic and socio-demographic features characterizing each pyrome and regime, we project them into a two-dimensional map to identify relevant patterns. Plotting and analyzing their demographic and climatic

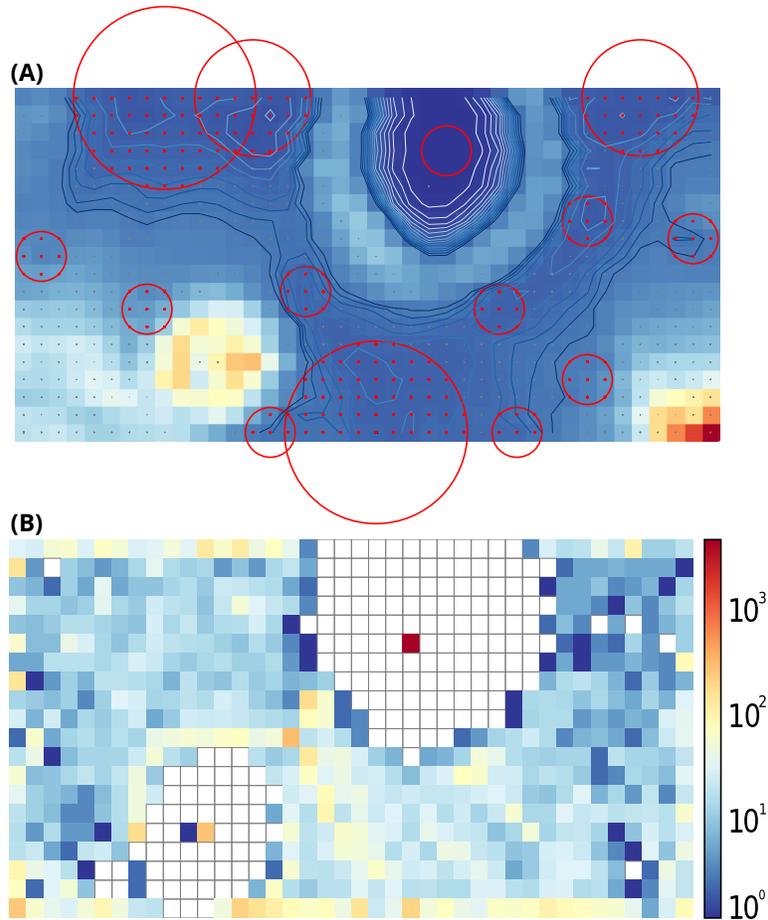


Figure 7.2: Discovering pyromes. Self-organizing maps are useful for summarizing multidimensional fire data and for determining potential groups of similar characteristics. These data are reduced to a two-dimensional grid and samples are organized according to their Euclidean distance. Observations sharing similar characteristics are easily visualized in a topographic map (A) where warmer colors represent widely separated samples and cooler colors depict closely related values. Using image processing algorithms (see Section 7.3), we detect significant potential pyromes/clusters (red circles). The number of observations belonging to each section of the map can be presented in a matrix known as hit-map (B). As an example, we can easily observe the group of cells without fire activity as a large dark blue region (top) and white valley (bottom), representing a significant percentage of the observations.

components (Figure 7.5) we observe: a) Regimes associated with low fire activity pyromes (R1, R8, and R9) show similar dispersion and trends for their demographic and climatic components. These regimes are commonly located in dry, cold, and sparsely populated ar-

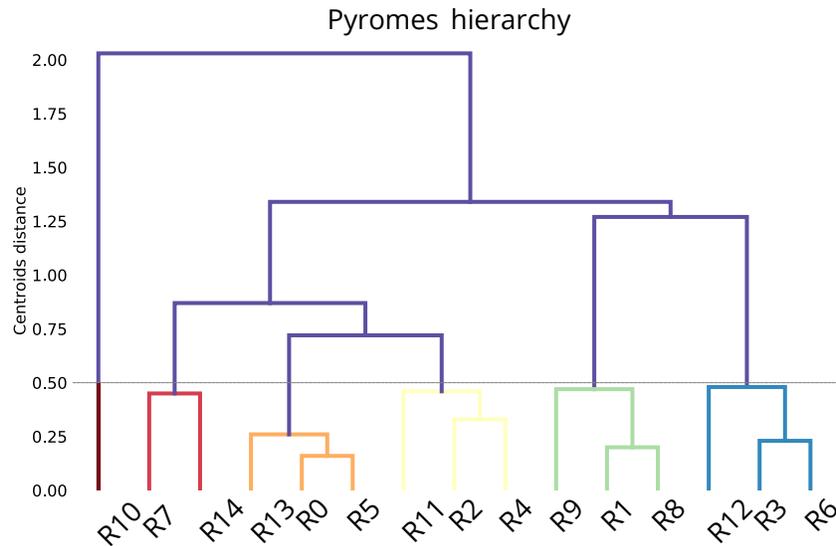


Figure 7.3: Pyromes’ hierarchy. Dendrogram summarizing the hierarchy of the determined fifteen pyromes and similarities between them. As observed, the pyromes can be collapsed into six macro groups sharing fire behavior characteristics, consistent with our statistical results (see Appendix E), where, for example, the pyrome of observations representing extreme and rare events (R10) is clearly independent of other clusters.

eas, with an exception for the regimes associated with the R8 pyrome. These regimes often appear in more densely populated regions of the world and under less extreme weather conditions. In all cases, no common land uses were found to define fire activity. No pattern is observed, with land covers varying from a predominance of evergreen forest in warmer areas, to grasslands and sparse vegetation in colder regions, whereas in milder and more populated areas, there is often the presence of mosaics of different land uses; b) Regimes experiencing large fires (within pyromes R2, R4, R10, and R11) are often located on balanced areas where no predominance on any of the demographic or climatic components can be observed, but focusing on colder regions for regimes within the R4 pyrome, the one with the milder fire activity of the group. The most common land uses within the area of these regimes are grasslands, savannas, and shrublands. c) Regimes belonging to pyromes where large but slow spread fires are present (R4 and R7), are always framed by cold, dry, and non-populated areas, covered mainly by savannas and shrublands. In this case, regimes and the associated pyromes share very similar climatic and socio-economic characteristics; d) Regimes where small but very frequent fires are characteristic (pyromes R3, R6, and R12) are located in densely populated areas, with less than average incomes, where warm and moist conditions plus the dominance of savannas and shrublands are common. However, a singularity can be observed in South-Asia, where the observed fire conditions are related

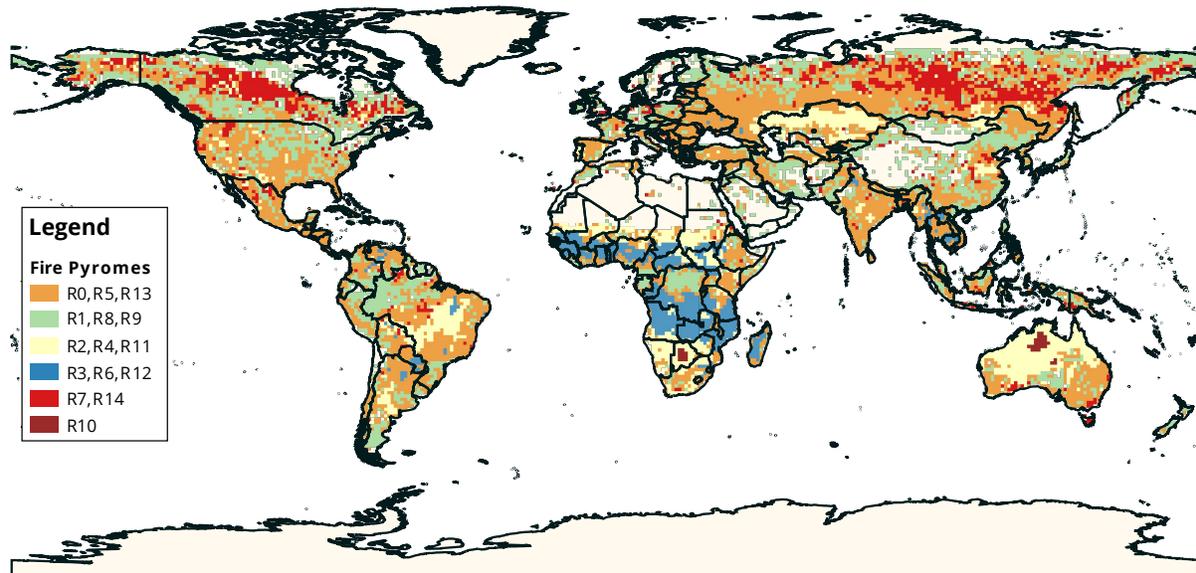
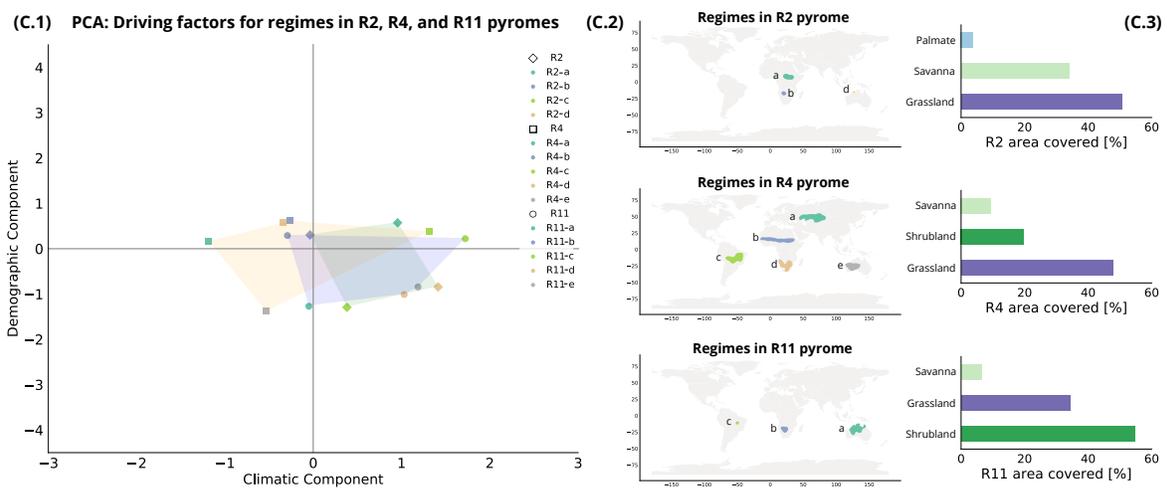
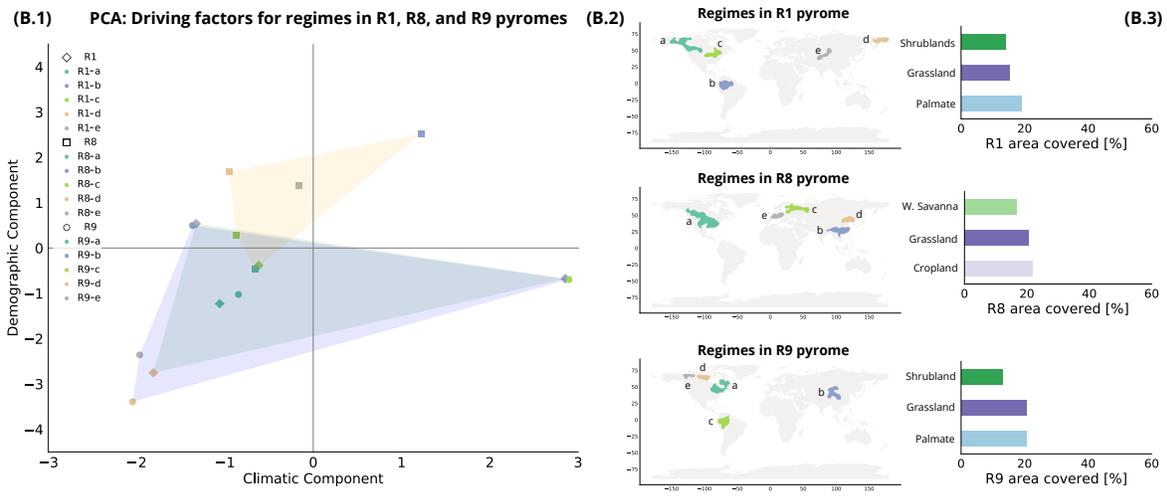
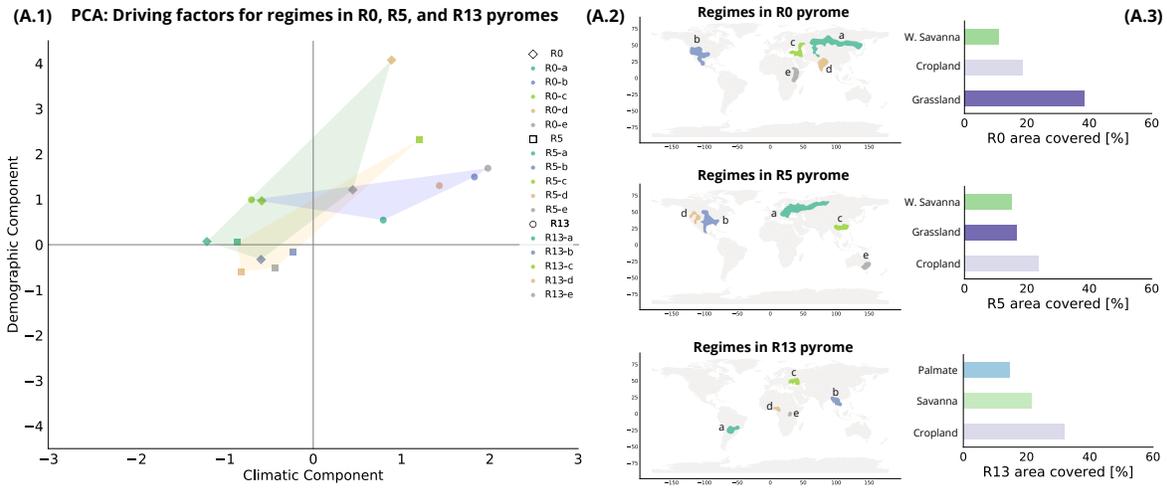


Figure 7.4: Fire pyromes. Cells covering the gridded world are classified into the six macro groups determined in the dendrogram (different colors) composed by the determined pyromes based on historical averages of fire characteristics including average annual burning frequency, size, perimeter, duration, daily expansion, and perimeter to area ratio values. The white and light brown backgrounds represent sections of the world where no fire events were registered for the studied period. From the results, certain pyromes cover multiple regions of the globe that do not seem to be related by climatic or demographic conditions. In order to understand their composition, we need to determine and characterize those areas where the observations of each group form a compact and well-defined cluster (via a density analysis) identifying fire regimes associated with specific regions.

to more diverse landscapes in which evergreen forest, savannas, and croplands cohabitate; e) Regimes associated with medium-small size fires but far less frequent than the previous ones, tend to be present under a wider range of conditions, as observed in their associated pyromes (R0, R5, and R13). These regimes can occur in areas dominated by grasslands, croplands, and savannas, but we can also find large areas covered by mixed or evergreen forest, leading to a higher level of landscape fragmentation. In general, we note that low population densities are common for those regimes located in dry and cold regions. On the other hand, the density of the human population increases sharply when these regimes are located in warmer and moister regions.

Although an analysis of combined driving factors by regimes provides a deeper understanding of the variability of conditions associated with these regimes, it still captures only a portion of the full picture required to understand fire regimes. To accomplish this goal, the best way is to observe each of the regimes and the potential influencing factors indepen-



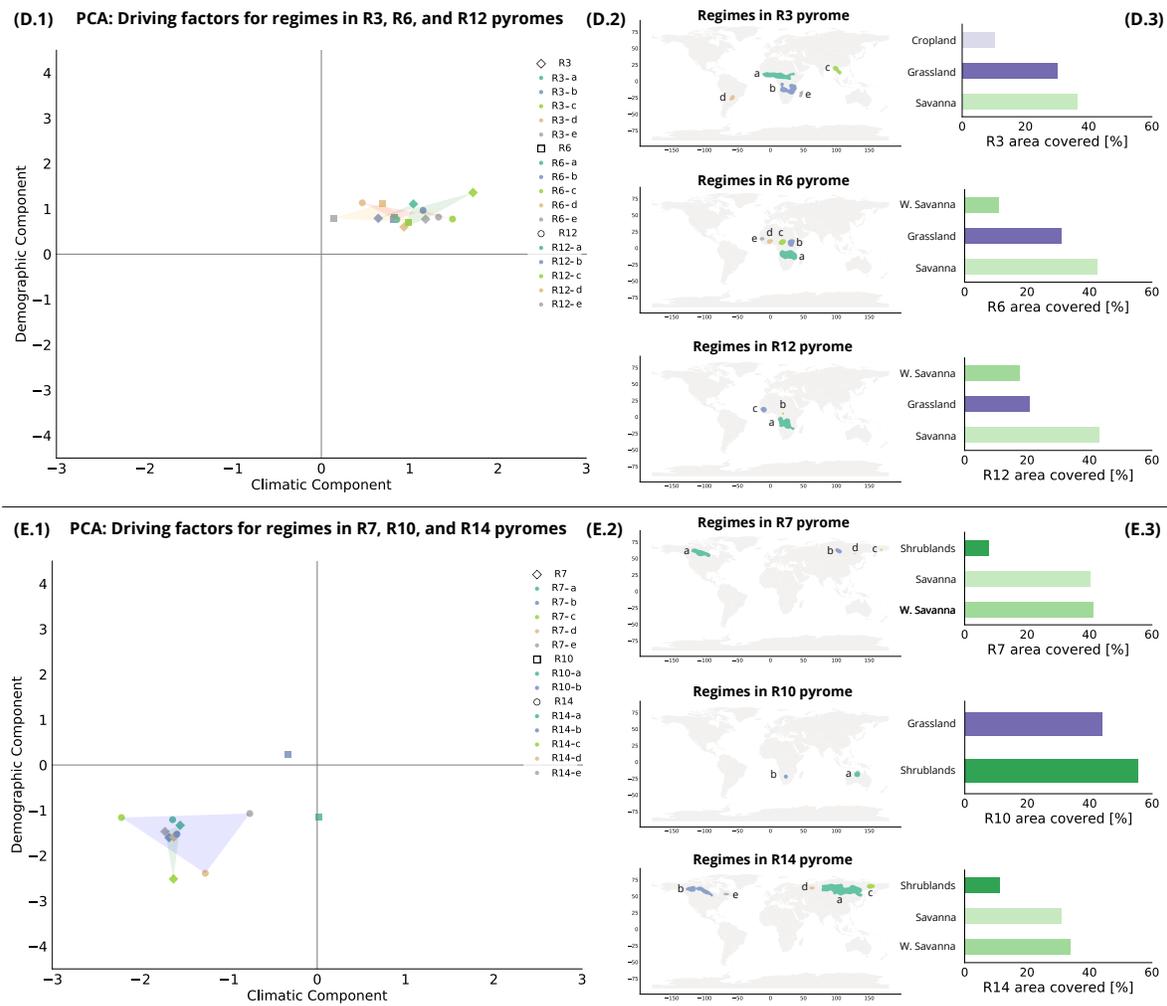


Figure 7.5: Regimes and driving factors. Demographic and climatic variables of the most relevant regimes per pyrome are summarized into two independent components using a principal component analysis (PCA) to visualize the regimes’ driving factors (A.1, B.1, C.1, D.1, and E.1 panels). A denser population, lower GDP, and harder accessibility characterize regimes located in higher values on the demographic component (y-axis). Regimes located in higher values on the climatic component experience higher average temperatures, higher precipitation levels, and lower PDSI values (x-axis). A shaded region connects regimes within the same pyrome, highlighting the dispersion of the regimes’ potential driving factors within each pyrome. The spatial location of the regimes (A.2, B.2, C.2, D.2, and E.2 panels, matching the colors of the regimes in the PCA plots) and the proportion of the dominant land covers per pyrome (A.3, B.3, C.3, D.3, and E.3 panels) are included for reference, providing an overall comprehensive characterization of the most relevant regimes (62 in total) for each pyrome. Detailed results per regime including their dominant land-covers can be found in Appendix E.

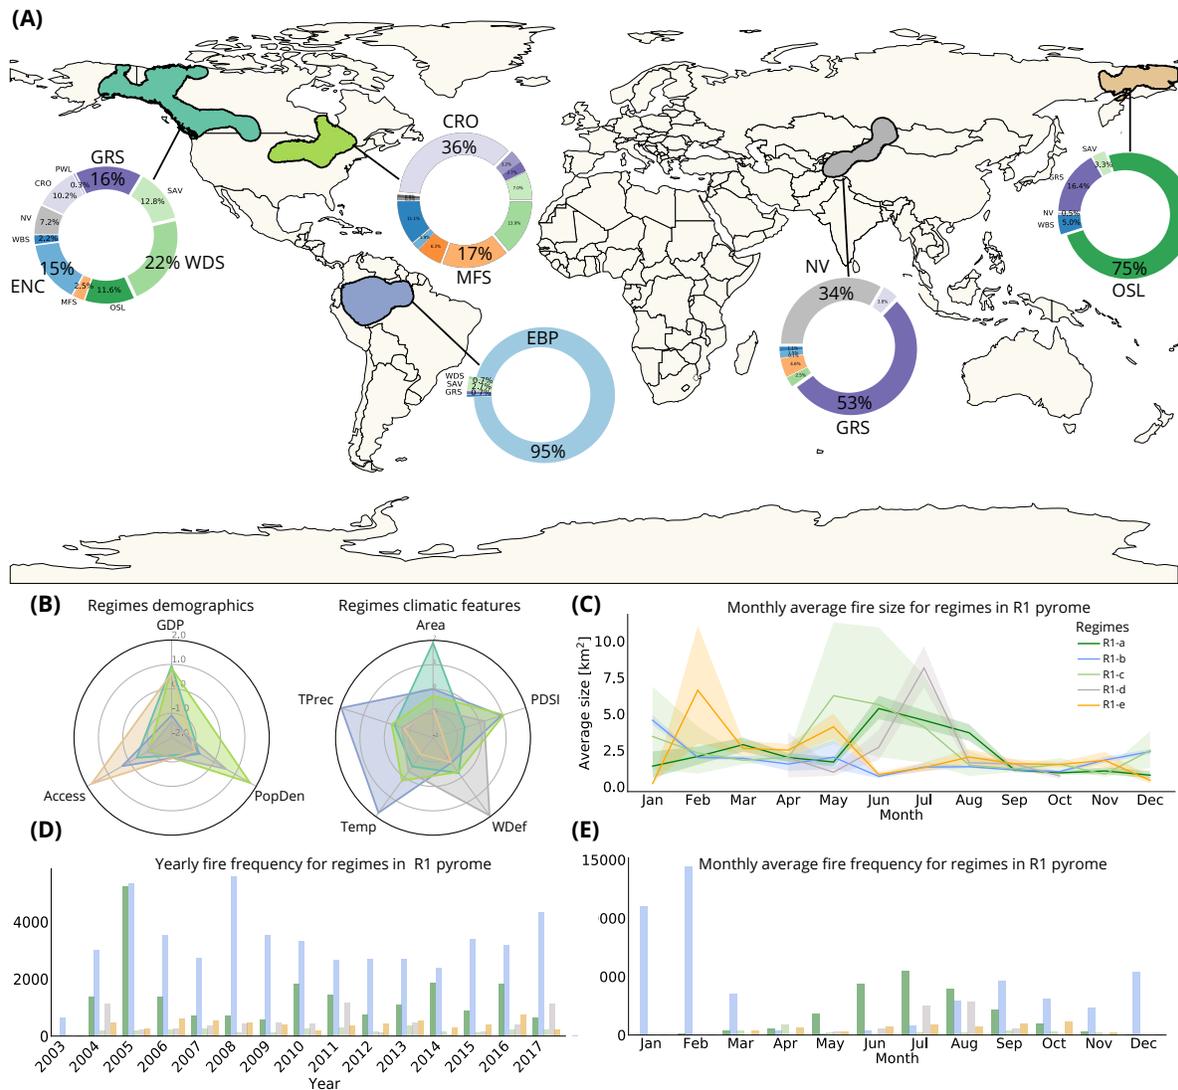


Figure 7.6: Characterizing fire regimes. Five largest spatial regimes are determined and represented with different colors for R1 pyrome after a kernel and contour level analysis. Despite belonging to the same pyrome, dense observation areas are spread in regions with very different climatic and demographic characteristics. In this pyrome, regimes cover parts of the western coast of Canada and Alaska (dark green), the Amazonian forest of Peru and Brazil (blue), the North America great lakes area (green), the eastern extreme of Russia (orange), and the central Asia region (gray). Significant differences can be observed between these regions in (A) location and land-cover distribution; (B) socio-demographic and climatic attributes; and (C-E) seasonal fire characteristics. Land covers observed in the pyrome include savannas (WDS, SAV), grasslands (GRS), conifers (ENC), evergreen broadleaf palmate (EBP), shrublands (OSL), mixed forest (MFS), water bodies (WBS), non-vegetated (NV), croplands (CRO), and permanent wetlands (PWL). Similar comparisons and deeper assessments can be performed for all regimes (see Appendix E).

dently. As an example, we explore the regimes within the R1 pyrome (Figure 7.6), where differences between regimes are especially clear between the Amazonian regime, denoted in blue, and the two regimes located on the northern hemisphere denoted by different shades of green. In this example, we found that the largest regime (R1-a) located in the western part of North America describes the characteristic of low fire activity (1.63 fires per year on average) of R1 regime, driven by cold temperatures and sparse vegetation on the inland parts of the area and probably an intense suppression on the coastal zones. On the other hand, the Amazonian hot-spot (R1-b) is influenced by very high moisture (PDSI of 75.39, 0.23 m of average yearly precipitation) due to rainfall and closed vegetation coverage. The hot-spot located in the area of large lakes (R1-c) is characterized by a heterogeneous landscape of mixed forests (16.8%) and croplands (35.9%), which together with suppression policies may justify its limited fire activity. From the two regimes located in Asia, the one on the eastern part of Siberia (R1-e, orange) is characterized by cold weather (an average of -8.88 C through the year), low population density (0.056 individuals per km^2), and sparse vegetation. On the other hand, the area in central Asia (R1-d) is clearly defined by the lack of water (average water deficit of 415.79 mm) and absence of large plants (86% of the land covered by a combination of grasslands and non-vegetated areas). Considering the omission of seasonality from the metrics used to define the pyromes, differences in this regard can be observed a posteriori among the regions. Whereas some subgroups show a clear tendency to aggregate the number of fires and area burned in summer in the northern hemisphere, the Amazonian regime shows a lesser tendency to aggregate over a specific season, with the peak fire activity between December and January, but still occurring in the rest of the year. The eastern Siberian regime also shows a different temporal pattern, having most of the fires in spring, but those of larger size during summer (see Appendix E for a comprehensive description of all pyromes and regimes).

The results show that different combinations of climatic, vegetation, and human factors may lead to similar fire patterns in different regions of the world. However, a change in those factors may induce swifts on fire activity to nearby locations [17, 315]. Nevertheless, it is also possible to determine clear gradients of fire activity if one of the influencing factors, especially climate or vegetation, changes accordingly. This spatial gradient is clear in those regimes with little fire activity and limited by low temperatures and scarce vegetation in northernmost boreal areas. As an example, we observe some regimes associated with pyrome R9 shifting into regimes characterized by the fire activity of R1, and, when fire-related conditions become more conducive, to regimes belonging to R7 or R14 pyromes. Similarly, regions with larger and fast expansive fires (e.g., characteristic of R10) are surrounded by pyromes sustaining smaller but still large fires (R11, and this one, surrounded by regimes within R4) if the climate and vegetation become gradually less hazardous. It is also possible to identify swifts on fire activity in nearby areas caused by non-natural processes. For example, in the limits of the Amazonian forest (R4-c) where continuous coverage of the rain forest has been disrupted [91, 92, 134] leading to sustain frequent fires of considerable size nowadays. Other natural processes can be observed across Asia, where slight differences in climate and subsequently in vegetation, e.g., increased rainfall in Kazakhstan (R4-a) or higher temperatures (R0-a),

modify the recurrence and size of fires. Contrary to this process of increased fire activity due to deforestation processes, in western Australia (R4-e), an arid environment is also able to sustain recurrent fires after stabilization of the allochthonous vegetation [67]. Therefore, it can be stated that there is a clear, but still complex relationship between the distribution of the driving factors and fire activity. Studies based on the impact changing factors have on fire activity, provide an important source of information about the past and future change of fires and their impact [49, 143, 176], but only an understanding of the interactions between factors in an area corresponding to a fire regime, and comparisons between regimes will provide the whole picture to understand why a regime is present on a certain region and not in another. Moreover, a similar approach should be conducted to predict possible shifts in regimes that, as has been shown can hardly be interpreted by the change of a unique factor.

The proposed methodology is able to spatially frame a fire regime on a specific area, with a high level of fire activity homogeneity, vegetation characteristics, and underlying vegetation factors. Constraints due to the different interpretations of fire regimes [191] have restrained this delimitation process at a global scale. This point is crucial, as it is recognized that understanding the fire regimes is the key to understand the present and future fire activity and its impact on the earth. The most limiting constraint for delimiting fire regimes at a global scale seems to be coupling the fire activity and the associated vegetation. This study did not assume that a single vegetation type has to be associated with a fire regime, as fire controls are heterogeneous across landscapes, and the spatial configuration of vegetation and fuels define the characteristics of a fire regime on a landscape [262, 321]. Assuming that, at a global scale, a mosaic of vegetation groups does not influence or can be part of a fire regime seems a very limiting premise. In this context, we believe that a global mosaic of vegetations types and other land covers, even if coarse, should be also considered as a part of what defines a regime, not a constraint. Moreover, if fire regimes are to be studied, they need to be framed spatially, and neither the results from defining equal fire pyromes nor the use of single vegetation groups can solve one of the most essential requirements for setting a fire regime: define the specific region where fire characteristics are repeated over time [191, 236].

7.6 Conclusions

Understanding and characterizing the wildfire regimes will provide crucial knowledge for a better understanding of both temporal and spatial impact of fire on the evolution of ecosystems, and on several services that are provided by them [65]. For example, perturbations in the regime have a significant impact on the whole ecosystem, particularly with vegetation [51]. A comprehension of the characteristics/trends and main features/factors that define the current regime will be crucial to support environmental conservation decisions in order to avoid shifts in undesirable and sometimes irreversible directions. This will allow national or supranational authorities to account for and whenever possible mitigate those changes supported by scientific and quantitative results/analysis.

The proposed framework and classification system allow the determination of fire regimes and their most common regions in a systematic way, without assuming that geopolitical borders or climatic characteristics of vegetation biomes are a constraint when framing their influencing area. Revealing similarities and dissimilarities between the fire pyromes and underlying factors provides the necessity of specific studies and adaptation measures across the globe, planned according to the relative importance of the driving factors and their expected change. This type of assessment, as shown by the study, requires a two-step clustering process. One based on fire characteristics alone and a second one focused on the spatial distribution of those fire characteristics. Without splitting a fire pyrome into spatially framed regimes, it has been clearly stated that understanding the underlying factors that cause such specific fire behavior becomes not only difficult but also may produce inconclusive or even misleading results. Maintaining, restoring, or mitigating fire activity depends on the understanding of current natural fire regimes, and their change under future conditions. This knowledge cannot be based on global generalizations but should base on site-specific recommendations. Site-specific considerations other than climate should be embedded into the models. Similarly, any global study related to fire activity or its future change should integrate the combination and interrelation of all influencing factors, often regionally specific, without assuming global oversimplifications on the impact of a single factor or its change. Delimiting fire regimes in the space, study them individually, compare distant regimes with similar fire activity, or adjacent fire regimes with differences in fire activity may be the key to truly understand the complexity of fire regimes.

Bibliography

- [1] John T Abatzoglou and A Park Williams. “Impact of anthropogenic climate change on wildfire across western US forests”. In: *Proceedings of the National Academy of Sciences* 113.42 (2016), pp. 11770–11775.
- [2] Mauricio A. Acuna et al. “Integrated spatial fire and forest management planning”. In: *Canadian Journal of Forest Research* 40.12 (2010), pp. 2370–2383. DOI: [10.1139/X10-151](https://doi.org/10.1139/X10-151). eprint: <https://doi.org/10.1139/X10-151>. URL: <https://doi.org/10.1139/X10-151>.
- [3] James K Agee and Carl N Skinner. “Basic principles of forest fuel reduction treatments”. In: *Forest ecology and management* 211.1-2 (2005), pp. 83–96.
- [4] Alan A Ager, Michelle A Day, and Kevin Vogler. “Production possibility frontiers and socioecological tradeoffs for restoration of fire adapted forests”. In: *Journal of Environmental Management* 176 (2016), pp. 157–168.
- [5] Alan A Ager, Nicole M Vaillant, and Mark A Finney. “A comparison of landscape fuel treatment strategies to mitigate wildland fire risk in the urban interface and preserve old forest structure”. In: *Forest Ecology and Management* 259.8 (2010), pp. 1556–1570.
- [6] Alan A Ager, Nicole M Vaillant, and Andrew McMahan. “Restoration of fire in managed forests: a model to prioritize landscapes and analyze tradeoffs”. In: *Ecosphere* 4.2 (2013), pp. 1–19.
- [7] Alan A Ager et al. “Modeling wildfire risk to northern spotted owl (*Strix occidentalis caurina*) habitat in Central Oregon, USA”. In: *Forest Ecology and Management* 246.1 (2007), pp. 45–56.
- [8] Stéphane Alarie et al. “Snow water equivalent estimation using blackbox optimization”. In: *Pac J Optim* 9.1 (2013), pp. 1–21.
- [9] Pedro Alberto et al. “Pattern search methods for user-provided points: Application to molecular geometry problems”. In: *SIAM Journal on Optimization* 14.4 (2004), pp. 1216–1236.
- [10] Frank A Albini. *Estimating wildfire behavior and effects*. Vol. 30. Department of Agriculture, Forest Service, Intermountain Forest and Range . . . , 1976.
- [11] Fermín J Alcasena et al. “Optimizing prescribed fire allocation for managing fire risk in central Catalonia”. In: *Science of the total environment* 621 (2018), pp. 872–885.

- [12] Ane A Alencar et al. “Landscape fragmentation, severe drought, and the new Amazon forest fire regime”. In: *Ecological applications* 25.6 (2015), pp. 1493–1505.
- [13] Martin E Alexander and BD Lawson. *User guide to the Canadian Forest Fire Behaviour Prediction System: rate of spread relationships*. Environment Canada, Canadian Forestry Service, 1984.
- [14] ME Alexander et al. “Estimating the length-to-breadth ratio of elliptical forest fire patterns”. In: *Proceedings of the eighth conference on fire and forest meteorology*. Vol. 29. Soc. Am. For Bethesda, MD. 1985, pp. 85–04.
- [15] A. Alexandridis et al. “A cellular automata model for forest fire spread prediction: The case of the wildfire that swept through Spetses Island in 1990”. In: *Applied Mathematics and Computation* 204.1 (2008), pp. 191–201.
- [16] Alex Alexandridis et al. “A cellular automata model for forest fire spread prediction: The case of the wildfire that swept through Spetses Island in 1990”. In: *Applied Mathematics and Computation* 204.1 (2008), pp. 191–201.
- [17] Craig D Allen et al. “Ecological restoration of southwestern ponderosa pine ecosystems: a broad perspective”. In: *Ecological applications* 12.5 (2002), pp. 1418–1433.
- [18] William M Alley. “The Palmer drought severity index: limitations and assumptions”. In: *Journal of climate and applied meteorology* 23.7 (1984), pp. 1100–1109.
- [19] Daniel Aloise et al. “NP-hardness of Euclidean sum-of-squares clustering”. In: *Machine learning* 75.2 (2009), pp. 245–248.
- [20] Adison Altamirano et al. “Influencia de la heterogeneidad del paisaje en la ocurrencia de incendios forestales en Chile Central”. In: *Revista de Geografía Norte Grande* 55 (2013), pp. 157–170.
- [21] Eduardo Álvarez-Miranda, Ivana Ljubić, and Petra Mutzel. “The maximum weight connected subgraph problem”. In: *Facets of Combinatorial Optimization*. Springer, 2013, pp. 245–270.
- [22] Giuseppe Amatulli et al. “Assessing long-term fire risk at local scale by means of decision tree technique”. In: *Journal of Geophysical Research: Biogeosciences* 111.G4 (2006).
- [23] Niels Andela et al. “A human-driven decline in global burned area”. In: *Science* 356.6345 (2017), pp. 1356–1362.
- [24] Niels Andela et al. “The Global Fire Atlas of individual fire size, duration, speed and direction”. In: *Earth System Science Data* 11.2 (2019), pp. 529–552.
- [25] D. H. Anderson et al. “Modelling the spread of grass fires”. In: *The ANZIAM Journal* 23.4 (1982), pp. 451–466.
- [26] H. E. Anderson. “Aids to determining fuel models for estimating fire behavior.” In: *US Dep. Agric., For. Serv., Intermt. For. Range Exp. Stn., Ogden, UT Gen Tech. Rep. INT 122* (1982), p. 22.

- [27] Mihael Ankerst et al. “OPTICS: ordering points to identify the clustering structure”. In: *ACM Sigmod record* 28.2 (1999), pp. 49–60.
- [28] B. Arca et al. “A web-based wildfire simulator for operational applications”. In: *International Journal of Wildland Fire* 28.2 (2019), pp. 99–112.
- [29] Sally Archibald, A Carla Staver, and Simon A Levin. “Evolution of human-driven fire regimes in Africa”. In: *Proceedings of the National Academy of Sciences* 109.3 (2012), pp. 847–852.
- [30] Sally Archibald et al. “Defining pyromes and global syndromes of fire regimes”. In: *Proceedings of the National Academy of Sciences* 110.16 (2013), pp. 6442–6447.
- [31] Juan Pablo Argañaraz et al. “Assessing wildfire exposure in the wildland-urban interface area of the mountains of central Argentina”. In: *Journal of environmental management* 196 (2017), pp. 499–510.
- [32] Tomàs Artés et al. “A global wildfire dataset for the analysis of fire regimes and fire behaviour”. In: *Scientific data* 6.1 (2019), pp. 1–11.
- [33] Charles Audet and Warren Hare. *Derivative-free and blackbox optimization*. Springer, 2017.
- [34] Charles Audet and Dominique Orban. “Finding optimal algorithmic parameters using derivative-free optimization”. In: *SIAM Journal on Optimization* 17.3 (2006), pp. 642–664.
- [35] Paolo Avesani, Anna Perini, and Francesco Ricci. “Interactive case-based planning for forest fire management”. In: *Applied Intelligence* 13.1 (2000), pp. 41–57.
- [36] Drew Bagnell and Andrew Y Ng. “On local rewards and scaling distributed reinforcement learning”. In: *Advances in Neural Information Processing Systems*. 2006, pp. 91–98.
- [37] Atoossa Bakhshaii and Edward A Johnson. “A review of a new generation of wildfire-atmosphere modeling”. In: *Canadian Journal of Forest Research* 49.6 (2019), pp. 565–574.
- [38] Erin Baldassari. “Camp Fire death toll grows to 29, matching 1933 blaze as state’s deadliest”. In: *East Bay Times (November 12, 2018)* (2018).
- [39] Kevin Barnett et al. “Beyond fuel treatment effectiveness: Characterizing Interactions between fire and treatments in the US”. In: *Forests* 7.10 (2016), p. 237.
- [40] Ana MG Barros et al. “Improving long-term fuel treatment effectiveness in the National Forest System through quantitative prioritization”. In: *Forest Ecology and Management* 433 (2019), pp. 514–527.
- [41] Peter Baumann et al. “Big data analytics for earth sciences: the EarthServer approach”. In: *International Journal of Digital Earth* 9.1 (2016), pp. 3–29.

- [42] Thomas Begin et al. “A DFO technique to calibrate queueing models”. In: *Computers & Operations Research* 37.2 (2010), pp. 273–281.
- [43] Akli Benali et al. “Bimodal fire regimes unveil a global-scale anthropogenic fingerprint”. In: *Global Ecology and Biogeography* 26.7 (2017), pp. 799–811.
- [44] Jennifer L Beverly and Neal McLoughlin. “Burn probability simulation and subsequent wildland fire activity in Alberta, Canada—implications for risk assessment and strategic planning”. In: *Forest Ecology and Management* 451 (2019), p. 117490.
- [45] Michael Bevers, Philip N Omi, and John Hof. “Random location of fuel treatments in wildland community interfaces: a percolation approach”. In: *Canadian Journal of Forest Research* 34.1 (2004), pp. 164–173.
- [46] Tanveer Hossain Bhuiyan et al. “A stochastic programming model with endogenous uncertainty for incentivizing fuel reduction treatment under uncertain landowner behavior”. In: *European Journal of Operational Research* 277.2 (2019), pp. 699–718.
- [47] Richard A Birdsey. *Carbon storage and accumulation in United States forest ecosystems*. Vol. 59. US Department of Agriculture, Forest Service, 1992.
- [48] Ekaba Bisong. “Google AutoML: Cloud Vision”. In: *Building Machine Learning and Deep Learning Models on Google Cloud Platform*. Springer, 2019, pp. 581–598.
- [49] I Bistinas et al. “Causal relationships vs. emergent patterns in the global controls of fire frequency”. In: *Biogeosciences* 11 (2014), pp. 5087–5101.
- [50] Ben Bond-Lamberty et al. “Fire as the dominant driver of central Canadian boreal forest carbon balance”. In: *Nature* 450.7166 (2007), pp. 89–92.
- [51] William J Bond and Jon E Keeley. “Fire as a global ‘herbivore’: the ecology and evolution of flammable ecosystems”. In: *Trends in ecology & evolution* 20.7 (2005), pp. 387–394.
- [52] Christopher Bone and Suzana Dragičević. “Simulation and validation of a reinforcement learning agent-based model for multi-stakeholder forest management”. In: *Computers, Environment and Urban Systems* 34.2 (2010), pp. 162–174.
- [53] Andrew J Booker et al. “Optimization using surrogate objectives on a helicopter test example”. In: *Computational Methods for Optimal Design and Control*. Springer, 1998, pp. 49–58.
- [54] Andrew Booker et al. “Managing surrogate objectives to optimize a helicopter rotor design—further experiments”. In: *7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. 1998, p. 4717.
- [55] Luigi Boschetti et al. “Global validation of the collection 6 MODIS burned area product”. In: *Remote Sensing of Environment* 235 (2019), p. 111490.
- [56] David MJS Bowman et al. “Human–environmental drivers and impacts of the globally extreme 2017 Chilean fires”. In: *Ambio* 48.4 (2019), pp. 350–362.

- [57] David MJS Bowman et al. “The human dimension of fire regimes on Earth”. In: *Journal of biogeography* 38.12 (2011), pp. 2223–2236.
- [58] David MJS Bowman and David MJS Bowman. *Australian rainforests: islands of green in a land of fire*. Cambridge University Press, 2000.
- [59] David MJS Bowman and Fay H Johnston. “Wildfire smoke, fire management, and human health”. In: *EcoHealth* 2.1 (2005), pp. 76–80.
- [60] David MJS Bowman, Jessica A O’Brien, and Johann G Goldammer. “Pyrogeography and the global quest for sustainable fire management”. In: *Annual Review of Environment and Resources* 38 (2013).
- [61] David MJS Bowman et al. “Vegetation fires in the Anthropocene”. In: *Nature Reviews Earth & Environment* 1.10 (2020), pp. 500–515.
- [62] D. Boychuk et al. “A stochastic fire growth model”. In: *Environmental and Ecological Statistics* 16 (2009), pp. 133–151.
- [63] Ulrik Brandes. “A faster algorithm for betweenness centrality”. In: *Journal of mathematical sociology* 25.2 (2001), pp. 163–177.
- [64] Ulrik Brandes. “A faster algorithm for betweenness centrality”. In: *Journal of mathematical sociology* 25.2 (2001), pp. 163–177.
- [65] Eckehard G Brockerhoff et al. *Forest biodiversity, ecosystem functioning and the provision of ecosystem services*. 2017.
- [66] Greg Brockman et al. “Openai gym”. In: *arXiv preprint arXiv:1606.01540* (2016).
- [67] Matthew L Brooks et al. “Effects of invasive alien plants on fire regimes”. In: *BioScience* 54.7 (2004), pp. 677–688.
- [68] Marcel Buchhorn et al. “Copernicus global land cover layers—collection 2”. In: *Remote Sensing* 12.6 (2020), p. 1044.
- [69] Robert E Burgan. *Behave: fire behavior prediction and fuel modeling system, fuel subsystem*. Vol. 167. Intermountain Forest and Range Experiment Station, Forest Service, US . . . , 1984.
- [70] George M Byram. “Combustion of forest fuels”. In: *Forest fire: control and use* (1959), pp. 61–89.
- [71] David E Calkin et al. “A comparative risk assessment framework for wildland fire management: the 2010 cohesive strategy science report”. In: (2011).
- [72] David E Calkin et al. “How risk management can prevent future wildfire disasters in the wildland-urban interface”. In: *Proceedings of the National Academy of Sciences* 111.2 (2014), pp. 746–751.
- [73] Maria Calviño-Cancela et al. “Wildfire risk associated with different vegetation types within and outside wildland-urban interfaces”. In: *Forest Ecology and Management* 372 (2016), pp. 1–9.

- [74] Yongcan Cao et al. “An overview of recent progress in the study of distributed multi-agent coordination”. In: *IEEE Transactions on Industrial informatics* 9.1 (2012), pp. 427–438.
- [75] Rodolfo Carvajal et al. “Imposing connectivity constraints in forest planning models”. In: *Operations Research* 61.4 (2013), pp. 824–836.
- [76] Wayne E Cascio. “Wildland fire smoke and human health”. In: *Science of the total environment* 624 (2018), pp. 586–595.
- [77] Sarah Cascone. *The Getty Museum Closes Its Doors as Wildfires Rage in California, But Says the Art Is Safe Inside Its Fireproof Building*. Oct. 2019. URL: <https://news.artnet.com/art-world/getty-museum-closed-wildfires-1691749>.
- [78] Davide Castelvechi. “Can we open the black box of AI?” In: *Nature News* 538.7623 (2016), p. 20.
- [79] EA Catchpole, TJ Hatton, and WR Catchpole. “Fire spread through nonhomogeneous fuel modelled as a Markov process”. In: *Ecological Modelling* 48.1-2 (1989), pp. 101–112.
- [80] Bureau Census. “U.S. Census Bureau QuickFacts: Napa County, California”. In: *Census Bureau QuickFacts* (2019). URL: <https://www.census.gov/quickfacts/napacountycalifornia>.
- [81] Yu-Han Chang, Tracey Ho, and Leslie P Kaelbling. “All learning is local: Multi-agent learning in global reward games”. In: *Advances in neural information processing systems*. 2004, pp. 807–814.
- [82] Nitesh V Chawla et al. “SMOTE: synthetic minority over-sampling technique”. In: *Journal of artificial intelligence research* 16 (2002), pp. 321–357.
- [83] Keping Chen and John McAneney. “Quantifying bushfire penetration into urban areas in Australia”. In: *Geophysical Research Letters* 31.12 (2004).
- [84] Kyunghyun Cho et al. “On the properties of neural machine translation: Encoder-decoder approaches”. In: *arXiv preprint arXiv:1409.1259* (2014).
- [85] François Chollet. “Xception: Deep learning with depthwise separable convolutions”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 1251–1258.
- [86] Sylvain Christin, Eric Hervet, and Nicolas Lecomte. “Applications for deep learning in ecology”. In: *Methods in Ecology and Evolution* 10.10 (2019), pp. 1632–1644.
- [87] Woodam Chung. “Optimizing fuel treatments to reduce wildland fire risk”. In: *Current Forestry Reports* 1.1 (2015), pp. 44–51.
- [88] Emilio Chuvieco et al. “Historical background and current developments for mapping burned area from satellite Earth observation”. In: *Remote sensing of environment* 225 (2019), pp. 45–64.

- [89] Emilio Chuvieco, Louis Giglio, and Chris Justice. “Global characterization of fire activity: toward defining fire regimes from Earth observation data”. In: *Global change biology* 14.7 (2008), pp. 1488–1502.
- [90] Antonio Ciampi and Yves Lechevallier. “Clustering large, multi-level data sets: an approach based on kohonen self organizing maps”. In: *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer. 2000, pp. 353–358.
- [91] Mark A Cochrane. “Fire science for rainforests”. In: *Nature* 421.6926 (2003), pp. 913–919.
- [92] Mark A Cochrane and Mark D Schulze. “Fire as a Recurrent Event in Tropical Forests of the Eastern Amazon: Effects on Forest Structure, Biomass, and Species Composition 1”. In: *Biotropica* 31.1 (1999), pp. 2–16.
- [93] Mark A Cochrane et al. “Positive feedbacks in the fire dynamic of closed canopy tropical forests”. In: *Science* 284.5421 (1999), pp. 1832–1835.
- [94] Andrew R Conn, Katya Scheinberg, and Ph L Toint. “Recent progress in unconstrained nonlinear optimization without derivatives”. In: *Mathematical programming* 79.1-3 (1997), p. 397.
- [95] Andrew R Conn, Katya Scheinberg, and Luis N Vicente. *Introduction to derivative-free optimization*. Vol. 8. Siam, 2009.
- [96] Andrew R Conn and Philippe L Toint. “An algorithm using quadratic interpolation for unconstrained derivative free optimization”. In: *Nonlinear optimization and applications*. Springer, 1996, pp. 27–47.
- [97] Jon M Conrad et al. “Wildlife corridors as a connected subgraph problem”. In: *Journal of Environmental Economics and Management* 63.1 (2012), pp. 1–18.
- [98] Jennifer K Costanza et al. *Describing and analyzing landscape patterns: where are we now, and where are we going?* 2019.
- [99] Wildland Fire Leadership Council. “A national cohesive wildland fire management strategy”. In: *US Forest Service* (2014).
- [100] W Wallace Covington et al. “Restoring ecosystem health in ponderosa pine forests of the Southwest”. In: *Journal of Forestry, Vol. 95, No. 4* (1997).
- [101] Antonia Creswell et al. “Generative adversarial networks: An overview”. In: *IEEE Signal Processing Magazine* 35.1 (2018), pp. 53–65.
- [102] Thomas Curt, Thibaut Fréjaville, and Sébastien Lahaye. “Modelling the spatial patterns of ignition causes and fire regime features in southern France: implications for fire prevention policy”. In: *International Journal of Wildland Fire* 25.7 (2016), pp. 785–796.
- [103] L. Dagum and R. Menon. “OpenMP: an industry standard API for shared-memory programming”. In: *IEEE computational science and engineering* 5.1 (1998), pp. 46–55.

- [104] Andrea De Montis et al. “Landscape fragmentation in Mediterranean Europe: A comparative approach”. In: *Land use policy* 64 (2017), pp. 83–94.
- [105] *DeepMind AI reduces google data centre cooling bill*. <https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>. Accessed: 2021-02-11.
- [106] Scott J DeWald, Scott John Josiah, and Rebecca Erdkamp. *Heating with wood: Producing, harvesting and processing firewood*. Cooperative Extension, Institute of Agriculture and Natural Resources . . . , 2005.
- [107] SH Doerr. “Santm. C”. In: *Global trends in wildfire and its impacts: perceptions versus realities in a changing world, Phil. Trans. R. Soc. B* 371 (), p. 1696.
- [108] Edward Dwyer et al. “Characterization of the spatio-temporal patterns of global fire activity using satellite imagery for the period April 1992 to March 1993”. In: *Journal of Biogeography* 27.1 (2000), pp. 57–69.
- [109] Edward Dwyer et al. “Global spatial and temporal distribution of vegetation fire as determined from satellite observations”. In: *International Journal of Remote Sensing* 21.6-7 (2000), pp. 1289–1302.
- [110] William P Edwards. “The New Normal: Living with Wildland Fire”. In: *Natural Resources & Environment* 33.3 (2019), pp. 30–33.
- [111] Donald A Falk. “Process-centred restoration in a fire-adapted ponderosa pine forest”. In: *Journal for Nature Conservation* 14.3-4 (2006), pp. 140–151.
- [112] M. A. Finney. “FARSITE: fire area simulator — model development and evaluation”. In: *Rev. ed. US Dept Agric., ForServ. Res. Pap. RMRS-RP-4* (2004), p. 48.
- [113] M. A. Finney. “The challenge of quantitative risk analysis for wildland fire”. In: *Forest Ecology and Management* 211 (2005), pp. 97–108.
- [114] Mark A Finney. “An overview of FlamMap fire modeling capabilities”. In: *In: Andrews, Patricia L.; Butler, Bret W., comps. 2006. Fuels Management-How to Measure Success: Conference Proceedings. 28-30 March 2006; Portland, OR. Proceedings RMRS-P-41. Fort Collins, CO: US Department of Agriculture, Forest Service, Rocky Mountain Research Station. p. 213-220. Vol. 41. 2006.*
- [115] Mark A Finney. “Calculation of fire spread rates across random landscapes”. In: *International Journal of Wildland Fire* 12.2 (2003), pp. 167–174.
- [116] Mark A Finney. “Design of regular landscape fuel treatment patterns for modifying fire growth and behavior”. In: *Forest Science* 47.2 (2001), pp. 219–228.
- [117] Mark A Finney. *FARSITE, Fire Area Simulator—model development and evaluation*. 4. US Department of Agriculture, Forest Service, Rocky Mountain Research Station, 1998.
- [118] Mark A Finney. “Fire growth using minimum travel time methods”. In: *Canadian Journal of Forest Research* 32.8 (2002), pp. 1420–1424.

- [119] Mark A Finney et al. “A method for ensemble wildland fire simulation”. In: *Environmental Modeling & Assessment* 16.2 (2011), pp. 153–167.
- [120] Mark A Finney et al. “A simulation of probabilistic wildfire risk components for the continental United States”. In: *Stochastic Environmental Research and Risk Assessment* 25.7 (2011), pp. 973–1000.
- [121] Mike D Flannigan et al. “Implications of changing climate for global wildland fire”. In: *International journal of wildland fire* 18.5 (2009), pp. 483–507.
- [122] Mike Flannigan, Brian Stocks, and Mike Weber. “Fire regimes and climatic change in Canadian forests”. In: *Fire and climatic change in temperate ecosystems of the western Americas*. Springer, 2003, pp. 97–119.
- [123] Jakob N Foerster et al. “Counterfactual multi-agent policy gradients”. In: *Thirty-second AAAI conference on artificial intelligence*. 2018.
- [124] Jakob Foerster et al. “Learning to communicate with deep multi-agent reinforcement learning”. In: *Advances in neural information processing systems*. 2016, pp. 2137–2145.
- [125] Forestry-Canada. *Development and structure of the Canadian Forest Fire Behavior Prediction System*. Tech. rep. Forestry Canada, Headquarters, Fire Danger Group, Science, and Sustainable Development Directorate, Ottawa. Information Report ST-X-3. 64 p., 1992.
- [126] E. Forgy. “Cluster analysis of multivariate data: efficiency versus interpretability of classifications”. In: *Biometrics* 21 (1965), pp. 768–780.
- [127] Félix-Antoine Fortin et al. “DEAP: Evolutionary algorithms made easy”. In: *Journal of Machine Learning Research* 13.Jul (2012), pp. 2171–2175.
- [128] Shodor Education Foundation. “How to weight a tree”. In: *Project SUCCEED-HI: For Students!* (2002). URL: <http://www.shodor.org/succeedhi/succeedhi/weightree/teacher/activities.html>.
- [129] Francis M Fujioka. “Estimating wildland fire rate of spread in a spatially nonuniform environment”. In: *Forest Science* 31.1 (1985), pp. 21–29.
- [130] Francis M Fujioka et al. “Fire danger and fire behavior modeling systems in Australia, Europe, and North America”. In: *Developments in Environmental Science* 8 (2008), pp. 471–497.
- [131] Thomas Fuller. *Three Weeks After Fire, Official Search for Dead Is Completed*. Nov. 2018. URL: <https://www.nytimes.com/2018/11/29/us/victims-california-fires-missing.html>.
- [132] Anne Ganteaume and Alexandra D Syphard. “Ignition Sources”. In: (2018).
- [133] Christina Geller. “Automated burned area identification in real-time during wildfire events using WorldView imagery for the insurance industry”. In: *Earth Resources and Environmental Remote Sensing/GIS Applications IX*. Vol. 10790. International Society for Optics and Photonics. 2018, p. 1079015.

- [134] Jeffrey J Gerwing. “Degradation of forests through logging and fire in the eastern Brazilian Amazon”. In: *Forest ecology and management* 157.1-3 (2002), pp. 131–141.
- [135] T. Ghisu et al. “An optimal cellular automata algorithm for simulating wildfire spread”. In: *Environmental Modelling & Software* 71 (2015), pp. 1–14.
- [136] Omid Ghorbanzadeh et al. “Spatial prediction of wildfire susceptibility using field survey gps data and machine learning approaches”. In: *Fire* 2.3 (2019), p. 43.
- [137] Louis Giglio, James T Randerson, and Guido R van der Werf. “Analysis of daily, monthly, and annual burned area using the fourth-generation global fire emissions database (GFED4)”. In: *Journal of Geophysical Research: Biogeosciences* 118.1 (2013), pp. 317–328.
- [138] Louis Giglio et al. “An enhanced contextual fire detection algorithm for MODIS”. In: *Remote sensing of environment* 87.2-3 (2003), pp. 273–282.
- [139] Louis Giglio et al. “The Collection 6 MODIS burned area mapping algorithm and product”. In: *Remote sensing of environment* 217 (2018), pp. 72–85.
- [140] Susana Gómez-González, Fernando Ojeda, and Paulo M Fernandes. “Portugal and Chile: Longing for sustainable forestry while rising from the ashes”. In: *Environmental Science & Policy* 81 (2018), pp. 104–107.
- [141] Susana Gómez-González et al. “Temperature and agriculture are largely associated with fire activity in Central Chile across different temporal periods”. In: *Forest Ecology and Management* 433 (2019), pp. 535–543.
- [142] Mauro E González et al. “The 2010–2015 Megadrought and its influence on the fire regime in central and south-central Chile”. In: *Ecosphere* 9.8 (2018), e02300.
- [143] Patrick Gonzalez et al. “Global patterns in the vulnerability of ecosystems to vegetation shifts due to climate change”. In: *Global Ecology and Biogeography* 19.6 (2010), pp. 755–768.
- [144] Noel Gorelick et al. “Google Earth Engine: Planetary-scale geospatial analysis for everyone”. In: *Remote sensing of Environment* 202 (2017), pp. 18–27.
- [145] Miranda E Gray and Brett G Dickson. “Applying fire connectivity and centrality measures to mitigate the cheatgrass-fire cycle in the arid West, USA”. In: *Landscape ecology* 31.8 (2016), pp. 1681–1696.
- [146] Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. “Variance reduction techniques for gradient estimates in reinforcement learning”. In: *Journal of Machine Learning Research* 5.Nov (2004), pp. 1471–1530.
- [147] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. “Cooperative multi-agent control using deep reinforcement learning”. In: *International Conference on Autonomous Agents and Multiagent Systems*. Springer. 2017, pp. 66–83.

- [148] Eric J Gustafson. “How has the state-of-the-art for quantification of landscape pattern advanced in the twenty-first century?” In: *Landscape Ecology* 34.9 (2019), pp. 2065–2072.
- [149] Aric Hagberg, Pieter Swart, and Daniel S Chult. *Exploring network structure, dynamics, and function using NetworkX*. Tech. rep. Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- [150] Aric Hagberg et al. “Networkx. High productivity software for complex networks”. In: *Webová stránka <https://networkx.lanl.gov/wiki>* (2013).
- [151] Stephanie E Hampton et al. “Big data and the future of ecology”. In: *Frontiers in Ecology and the Environment* 11.3 (2013), pp. 156–162.
- [152] Stijn Hantson et al. “Anthropogenic effects on global mean fire size”. In: *International Journal of Wildland Fire* 24.5 (2015), pp. 589–596.
- [153] WL Hare. “Using derivative free optimization for constrained parameter selection in a home and community care forecasting model”. In: *International perspectives on operations research and health care, Proceedings of the 34th meeting of the EURO working group on operational research applied to health sciences*. 2010, pp. 61–73.
- [154] William E Hart et al. *Pyomo-optimization modeling in python*. Vol. 67. Springer, 2017.
- [155] Haibo He and Eduardo A Garcia. “Learning from imbalanced data”. In: *IEEE Transactions on knowledge and data engineering* 21.9 (2009), pp. 1263–1284.
- [156] Kelvin G Hirsch et al. *Canadian forest fire behavior prediction (FBP) system: user’s guide*. Vol. 7. 1996.
- [157] Kelvin Hirsch et al. “Fire-smart forest management: A pragmatic approach to sustainable forest management in fire-dominated ecosystems”. In: *The Forestry Chronicle* 77.2 (2001), pp. 357–363.
- [158] Kelvin Hirsch et al. “Fire-smart forest management: a pragmatic approach to sustainable forest management in fire-dominated ecosystems”. In: *The Forestry Chronicle* 77.2 (2001), pp. 357–363.
- [159] Kelvin Hirsch et al. “Fire-smart forest management: a pragmatic approach to sustainable forest management in fire-dominated ecosystems”. In: *The Forestry Chronicle* 77.2 (2001), pp. 357–363.
- [160] Toan Minh Hoang et al. “Deep retinanet-based detection and classification of road markings by visible light camera sensors”. In: *Sensors* 19.2 (2019), p. 281.
- [161] Jonathan L Hodges and Brian Y Lattimer. “Wildland Fire Spread Modeling Using Convolutional Neural Networks”. In: *Fire Technology* 55.6 (2019), pp. 2115–2142.
- [162] kristina Houck and Patch Staff. *Top 20 most destructive California wildfires*. 2018. URL: <https://patch.com/california/san-diego/top-20-deadliest-most-destructive-wildfires-ca-history>.

- [163] Chen Huang et al. “Learning Deep Representation for Imbalanced Classification”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016.
- [164] Sergey Ioffe and Christian Szegedy. “Batch normalization: Accelerating deep network training by reducing internal covariate shift”. In: *arXiv preprint arXiv:1502.03167* (2015).
- [165] Piyush Jain et al. “A review of machine learning applications in wildfire science and management”. In: *arXiv preprint arXiv:2003.00646* (2020).
- [166] David S Johnson. “The NP-completeness column: an ongoing guide”. In: *Journal of Algorithms* 6.3 (1985), pp. 434–451.
- [167] Steven G Johnson. *The NLOpt nonlinear-optimization package*. 2014.
- [168] W Matt Jolly et al. “Climate-induced variations in global wildfire danger from 1979 to 2013”. In: *Nature communications* 6.1 (2015), pp. 1–11.
- [169] Emilio Jorge et al. “Learning to play guess who? and inventing a grounded language as a consequence”. In: *arXiv preprint arXiv:1611.03218* (2016).
- [170] G Julio, R Aguilera, and P Pedernera. “The Kitral System”. In: *Proc. International Workshop on Strategic Fire Planning Systems. USDA Forest Service, Fire Research Lab., Riverside, California (100 pp.)* 1997.
- [171] Mohannad Kabli, Jianbang Gan, and Lewis Ntaimo. “A stochastic programming model for fuel treatment management”. In: *Forests* 6.6 (2015), pp. 2148–2162.
- [172] Ammar Ismael Kadhim. “Survey on supervised machine learning techniques for automatic text classification”. In: *Artificial Intelligence Review* 52.1 (2019), pp. 273–292.
- [173] Eric S Kasischke and Merritt R Turetsky. “Recent changes in the fire regime across the North American boreal region—Spatial and temporal patterns of burning across Canada and Alaska”. In: *Geophysical research letters* 33.9 (2006).
- [174] Jon E Keeley and Alexandra D Syphard. “Twenty-first century California, USA, wildfires: fuel-dominated vs. wind-dominated fires”. In: *Fire Ecology* 15.1 (2019), p. 24.
- [175] Douglas I Kelley et al. “How contemporary bioclimatic and human controls change global fire regimes”. In: *Nature Climate Change* 9.9 (2019), pp. 690–696.
- [176] Luke T Kelly et al. “Fire and biodiversity in the Anthropocene”. In: *Science* 370.6519 (2020).
- [177] Joel K Kelso et al. “Techniques for evaluating wildfire simulators via the simulation of historical fires using the Australis simulator”. In: *International Journal of Wildland Fire* 24.6 (2015), pp. 784–797.
- [178] C. H. Key and N. C. Benson. “The Normalized Burn Ratio (NBR): A Landsat TM Radiometric Measure of Burn Severity, Northern Rocky Mountain Science Center. <http://www.nrmsc.usgs.gov/research/ndbr.htm>, 1999.” In: (1999).

- [179] Young-Hwan Kim, Pete Bettinger, and Mark Finney. “Spatial optimization of the pattern of fuel management activities and subsequent effects on simulated wildfires”. In: *European Journal of Operational Research* 197.1 (2009), pp. 253–265.
- [180] Heike Knicker. “How does fire affect the nature and stability of soil organic nitrogen and carbon? A review”. In: *Biogeochemistry* 85.1 (2007), pp. 91–118.
- [181] W Knorr, A Arneth, and L Jiang. “Demographic controls of future global fire risk”. In: *Nature Climate Change* 6.8 (2016), pp. 781–785.
- [182] Trupti M Kodinariya and Prashant R Makwana. “Review on determining number of Cluster in K-Means Clustering”. In: *International Journal* 1.6 (2013), pp. 90–95.
- [183] Walter D Koenig. “Spatial autocorrelation of ecological phenomena”. In: *Trends in Ecology & Evolution* 14.1 (1999), pp. 22–26.
- [184] Teuvo Kohonen. “Exploration of very large databases by self-organizing maps”. In: *Proceedings of international conference on neural networks (icnn'97)*. Vol. 1. IEEE. 1997, PL1–PL6.
- [185] Hui Kong, Hatice Cinar Akakin, and Sanjay E Sarma. “A generalized Laplacian of Gaussian filter for blob detection and its applications”. In: *IEEE transactions on cybernetics* 43.6 (2013), pp. 1719–1733.
- [186] P. H. Kourtz and W. G. O’Regan. “A model for a small forest fire . . . to simulate burned and burning areas for use in a detection model”. In: *For Sci* 17.2 (1971), pp. 163–169.
- [187] P. Kourtz, S. Nozaki, and W. G. O’Regan. “Forest fires in the computer — a model to predict the perimeter location of a forest fire.” In: *Fish Environ. Can Ottawa, ON. Inf. Rep. FF-X- 65* (1977).
- [188] Landon Kraemer and Bikramjit Banerjee. “Multi-agent reinforcement learning as a rehearsal for decentralized planning”. In: *Neurocomputing* 190 (2016), pp. 82–94.
- [189] H Anu Kramer et al. “Where wildfires destroy buildings in the US relative to the wildland–urban interface and national fire outreach programs”. In: *International journal of wildland fire* 27.5 (2018), pp. 329–341.
- [190] Meg A Krawchuk et al. “Global pyrogeography: the current and future distribution of wildfire”. In: *PloS one* 4.4 (2009).
- [191] Patrik Krebs et al. “Fire regime: history and definition of a key concept in disturbance ecology”. In: *Theory in Biosciences* 129.1 (2010), pp. 53–69.
- [192] Priya Krishnakumar. *We mapped every wine country fire. They’re larger and more destructive than ever.* Nov. 2019. URL: <https://www.latimes.com/projects/wine-country-fires-california-sonoma-napa/>.
- [193] Claudio A. Kuhlmann et al. “Generating Stochastic Ellipsoidal Forest and Wildland Fire Scar Scenarios for Strategic Forest Management Planning under Uncertainty”. In: *Forest Science* 61.3 (Dec. 2014), pp. 494–508.

- [194] Solomon Kullback and Richard A Leibler. “On information and sufficiency”. In: *The annals of mathematical statistics* 22.1 (1951), pp. 79–86.
- [195] Matthai Chakko Kuruvila and Ellen Lee. *Paradise fire evacuees starting to return home*. Feb. 2012. URL: <https://www.sfgate.com/bayarea/article/Paradise-fire-evacuees-starting-to-return-home-3209813.php>.
- [196] Paul S Laris. “Spatiotemporal problems with detecting and mapping mosaic fire regimes with coarse-resolution satellite data in savanna environments”. In: *Remote sensing of environment* 99.4 (2005), pp. 412–424.
- [197] Martin Lauer and Martin Riedmiller. “An algorithm for distributed reinforcement learning in cooperative multi-agent systems”. In: *In Proceedings of the Seventeenth International Conference on Machine Learning*. Citeseer. 2000.
- [198] William F Laurance and G Bruce Williamson. “Positive feedbacks among forest fragmentation, drought, and climate change in the Amazon”. In: *Conservation biology* 15.6 (2001), pp. 1529–1535.
- [199] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *nature* 521.7553 (2015), pp. 436–444.
- [200] Yann LeCun et al. “Backpropagation applied to handwritten zip code recognition”. In: *Neural computation* 1.4 (1989), pp. 541–551.
- [201] Tatsiana Levina et al. “Dynamic pricing with online learning and strategic consumers: an application of the aggregating algorithm”. In: *Operations Research* 57.2 (2009), pp. 327–341.
- [202] Sergey Levine et al. “End-to-end training of deep visuomotor policies”. In: *The Journal of Machine Learning Research* 17.1 (2016), pp. 1334–1373.
- [203] Yuxi Li. “Deep reinforcement learning: An overview”. In: *arXiv preprint arXiv:1701.07274* (2017).
- [204] Michael L Littman. “Markov games as a framework for multi-agent reinforcement learning”. In: *Machine learning proceedings 1994*. Elsevier, 1994, pp. 157–163.
- [205] Tianjia Liu et al. “Diagnosing spatial biases and uncertainties in global fire emissions inventories: Indonesia as regional case study”. In: *Remote Sensing of Environment* 237 (2020), p. 111557.
- [206] Weibo Liu et al. “A survey of deep neural network architectures and their applications”. In: *Neurocomputing* 234 (2017), pp. 11–26.
- [207] Zhihua Liu, Jian Yang, and Hong S He. “Studying the effects of fuel treatment based on burn probability on a boreal forest landscape”. In: *Journal of environmental management* 115 (2013), pp. 42–52.
- [208] Francisco Lloret et al. “Wildfires and landscape patterns in the Eastern Iberian Peninsula”. In: *Landscape Ecology* 17.8 (2002), pp. 745–759.

- [209] Craig Loehle. “Applying landscape principles to fire hazard reduction”. In: *Forest Ecology and management* 198.1-3 (2004), pp. 261–267.
- [210] David J Lohman, David Bickford, and Navjot S Sodhi. “The burning issue”. In: *Science* 316.5823 (2007), pp. 376–376.
- [211] Tengfei Long et al. “30 m Resolution Global Annual Burned Area Mapping Based on Landsat Images and Google Earth Engine”. In: *Remote Sensing* 11.5 (2019), p. 489.
- [212] Ryan Lowe et al. “Multi-agent actor-critic for mixed cooperative-competitive environments”. In: *Advances in neural information processing systems*. 2017, pp. 6379–6390.
- [213] Laurens van der Maaten and Geoffrey Hinton. “Visualizing data using t-SNE”. In: *Journal of machine learning research* 9.Nov (2008), pp. 2579–2605.
- [214] Donald W Marquardt. “An algorithm for least-squares estimation of nonlinear parameters”. In: *Journal of the society for Industrial and Applied Mathematics* 11.2 (1963), pp. 431–441.
- [215] Avi Bar Massada et al. “Wildfire ignition-distribution modelling: a comparative study in the Huron–Manistee National Forest, Michigan, USA”. In: *International journal of wildland fire* 22.2 (2013), pp. 174–183.
- [216] Laëtitia Matignon, Laurent Jeanpierre, and Abdel-illah Mouaddib. “Coordinated multi-robot exploration under communication constraints using decentralized markov decision processes”. In: *Twenty-sixth AAAI conference on artificial intelligence*. 2012.
- [217] Laëtitia Matignon, Guillaume J Laurent, and Nadine Le Fort-Piat. “Hysteretic q-learning: an algorithm for decentralized reinforcement learning in cooperative multi-agent teams”. In: *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2007, pp. 64–69.
- [218] Dmytro Matsypura, Oleg A Prokopyev, and Aizat Zahar. “Wildfire fuel management: network-based models and optimization of prescribed burning”. In: *European Journal of Operational Research* 264.2 (2018), pp. 774–796.
- [219] Alan Grant McArthur. *Weather and grassland fire behaviour*. Forestry and Timber Bureau, Department of national Development, Commonwealth . . . , 1966.
- [220] Kevin McGarigal. *FRAGSTATS: spatial pattern analysis program for quantifying landscape structure*. Vol. 351. US Department of Agriculture, Forest Service, Pacific Northwest Research Station, 1995.
- [221] Donald McKenzie, Carol Miller, and Donald A Falk. *The landscape ecology of fire*. Springer Science & Business Media, 2011.
- [222] Kendra K McLauchlan et al. “Fire as a fundamental ecological process: Research advances and frontiers”. In: *Journal of Ecology* 108.5 (2020), pp. 2047–2069.
- [223] David B McWethy et al. “Landscape drivers of recent fire activity (2001-2017) in south-central Chile”. In: *PLoS one* 13.8 (2018), e0201195.

- [224] Jay D Miller and Andrea E Thode. “Quantifying burn severity in a heterogeneous landscape with a relative version of the delta Normalized Burn Ratio (dNBR)”. In: *Remote Sensing of Environment* 109.1 (2007), pp. 66–80.
- [225] James Minas, John Hearne, and David Martell. “An integrated optimization model for fuel management and fire suppression preparedness planning”. In: *Annals of operations Research* 232.1 (2015), pp. 201–215.
- [226] Alejandro Miranda et al. “Evidence-based mapping of the wildland-urban interface to better identify human communities threatened by wildfires”. In: *Environmental Research Letters* (2020).
- [227] Alejandro Miranda et al. “Native forest loss in the Chilean biodiversity hotspot: revealing the evidence”. In: *Regional Environmental Change* 17.1 (2017), pp. 285–297.
- [228] Jason B Moats, Thomas J Chermack, and Larry M Dooley. “Using scenarios to develop crisis managers: Applications of scenario planning and scenario-based training”. In: *Advances in Developing Human Resources* 10.3 (2008), pp. 397–424.
- [229] Hossein Moayedi et al. “Fuzzy-metaheuristic ensembles for spatial assessment of forest fire susceptibility”. In: *Journal of environmental management* 260 (2020), p. 109867.
- [230] Jason J Moghaddas and Larry Craggs. “A fuel treatment reduces fire severity and increases suppression efficiency in a mixed conifer forest”. In: *International Journal of Wildland Fire* 16.6 (2008), pp. 673–678.
- [231] Cristina Montiel-Molina et al. “Have Historical Land Use/Land Cover Changes Triggered a Fire Regime Shift in Central Spain?” In: *Fire* 2.3 (2019), p. 44.
- [232] Francisco Moreira et al. “Landscape–wildfire interactions in southern Europe: implications for landscape management”. In: *Journal of environmental management* 92.10 (2011), pp. 2389–2402.
- [233] Francisco Moreira et al. “Wildfire management in Mediterranean-type regions: paradigm change needed”. In: *Environmental Research Letters* 15.1 (2020), p. 011001.
- [234] Francisco Moreira et al. “Size-dependent pattern of wildfire ignitions in Portugal: when do ignitions turn into big fires?” In: *Landscape Ecology* 25.9 (2010), pp. 1405–1417.
- [235] M Vanesa Moreno and Emilio Chuvieco. “Characterising fire regimes in Spain from fire statistics”. In: *International Journal of Wildland Fire* 22.3 (2013), pp. 296–305.
- [236] Penelope Morgan et al. “Mapping fire regimes across time and space: understanding coarse and fine-scale fire patterns”. In: *International Journal of Wildland Fire* 10.4 (2001), pp. 329–342.
- [237] Max A Moritz et al. “Learning to coexist with wildfire”. In: *Nature* 515.7525 (2014), p. 58.
- [238] Max A Moritz et al. “Climate change and disruptions to global fire activity”. In: *Ecosphere* 3.6 (2012), pp. 1–22.

- [239] Max A Moritz et al. “Wildfires, complexity, and highly optimized tolerance”. In: *Proceedings of the National Academy of Sciences* 102.50 (2005), pp. 17912–17917.
- [240] Pelagie Elimbi Moudio, Cristobal Pais, and Zuo-Jun Max Shen. “Quantifying the impact of ecosystem services for landscape management under wildfire hazard”. In: *Natural Hazards* 106.1 (2021), pp. 531–560.
- [241] Florent Mouillot and Christopher B Field. “Fire history and the global carbon budget: a 1×1 fire history reconstruction for the 20th century”. In: *Global Change Biology* 11.3 (2005), pp. 398–420.
- [242] Nicholas J Nauslar, John T Abatzoglou, and Patrick T Marsh. “The 2017 North Bay and Southern California fires: a case study”. In: *Fire* 1.1 (2018), p. 18.
- [243] John A Nelder and Roger Mead. “A simplex method for function minimization”. In: *The computer journal* 7.4 (1965), pp. 308–313.
- [244] Mohammad Sadegh Norouzzadeh et al. “Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning”. In: *Proceedings of the National Academy of Sciences* 115.25 (2018), E5716–E5725.
- [245] W. G. O’Regan, P. Kourtz, and S. Nozaki. “Bias in the contagion analog to fire spread”. In: *For Sci* 2.1 (1976), pp. 61–68.
- [246] W. G. O’Regan, S. Nozaki, and P. Kourtz. “A method for using directional rates of spread to predict forest fire configurations. Western States Section”. In: *The Combustion Institute WSS/CI Paper* 73-17 (1973).
- [247] Rodrigue Oeuvray. *Trust-region methods based on radial basis functions with application to biomedical imaging*. Tech. rep. EPFL, 2005.
- [248] Sandra Oliveira et al. “Modeling spatial patterns of fire occurrence in Mediterranean Europe using Multiple Regression and Random Forest”. In: *Forest Ecology and Management* 275 (2012), pp. 117–129.
- [249] Philip N Omi. “Theory and practice of wildland fuels management”. In: *Current Forestry Reports* 1.2 (2015), pp. 100–117.
- [250] Cristobal Pais et al. “Cell2Fire: A Cell Based Forest Fire Growth Model”. In: *arXiv preprint arXiv:1905.09317* (2019).
- [251] Cristobal Pais et al. “Cell2Fire: A Cell Based Forest Fire Growth Model”. In: *arXiv preprint arXiv:1905.09317* (2019).
- [252] Cristobal Pais et al. “Downstream protection value: Detecting critical zones for effective fuel-treatment under wildfire risk”. In: *Computers & Operations Research* 131 (2021), p. 105252.
- [253] Cristian D Palma et al. “Assessing the impact of stand-level harvests on the flammability of forest landscapes”. In: *International Journal of Wildland Fire* 16.5 (2007), pp. 584–592.

- [254] Liviu Panait and Sean Luke. “Cooperative multi-agent learning: The state of the art”. In: *Autonomous agents and multi-agent systems* 11.3 (2005), pp. 387–434.
- [255] G. D. Papadopoulos and F.-N. Pavlidou. “A comparative review on wild fire simulators”. In: *IEEE Syst. J.* 5 (2011), pp. 233–243.
- [256] George D Papadopoulos and Fotini-Niovi Pavlidou. “A comparative review on wildfire simulators”. In: *IEEE systems Journal* 5.2 (2011), pp. 233–243.
- [257] Marc-André Parisien et al. “Mapping wildfire susceptibility with the BURN-P3 simulation model”. In: *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Information report NOR-X-405.(Edmonton, AB)* (2005).
- [258] Marc-André Parisien, Dave R Junor, and Victor G Kafka. “Using Landscape-Based Decision Rules to Prioritize Locations of Fuel Treatments in the Boreal Mixedwood of Western Canada”. In: *In: Andrews, Patricia L.; Butler, Bret W., comps. 2006. Fuels Management-How to Measure Success: Conference Proceedings. 28-30 March 2006; Portland, OR. Proceedings RMRS-P-41. Fort Collins, CO: US Department of Agriculture, Forest Service, Rocky Mountain Research Station. p. 221-236.* Vol. 41. 2006.
- [259] Marc-André Parisien et al. “Applications of simulation-based burn probability modelling: a review”. In: *International journal of wildland fire* 28.12 (2020), pp. 913–926.
- [260] Marc-André Parisien et al. “Contributions of ignitions, fuels, and weather to the spatial patterns of burn probability of a boreal landscape”. In: *Ecosystems* 14.7 (2011), pp. 1141–1155.
- [261] MA Parisien et al. “Using the BURN-P3 simulation model to map wildfire susceptibility”. In: *Canadian Forest Service Report NOR-X-405* (2005).
- [262] Sean A Parks, Marc-André Parisien, and Carol Miller. “Spatial bottom-up controls on fire likelihood vary across western North America”. In: *Ecosphere* 3.1 (2012), pp. 1–20.
- [263] David J Parsons and Jan W van Wagtendonk. “Fire research and management in the Sierra Nevada National Parks”. In: *Science and ecosystem management in the national parks. University of Arizona Press, Tucson, Arizona, USA* (1996), pp. 25–48.
- [264] Juli G Pausas and Santiago Fernández-Muñoz. “Fire regime changes in the Western Mediterranean Basin: from fuel-limited to drought-driven fire regime”. In: *Climatic change* 110.1-2 (2012), pp. 215–226.
- [265] Patricio Pedernera and Guillermo Julio. “Improving the economic efficiency of combatting forest fires in Chile: the KITRAL system”. In: *Fire Economics, Planning, and Policy: Bottom Lines* (1999), p. 149.
- [266] Adam FA Pellegrini et al. “Fire frequency drives decadal changes in soil carbon and nitrogen and ecosystem productivity”. In: *Nature* 553.7687 (2018), pp. 194–198.
- [267] GLW Perry. “Current approaches to modelling the spread of wildland fire: a review”. In: *Progress in Physical Geography* 22.2 (1998), pp. 222–245.

- [268] Raechel A Portelli. “Don’t throw the baby out with the bathwater: reappreciating the dynamic relationship between humans, machines, and landscape images”. In: *Landscape Ecology* (2020), pp. 1–8.
- [269] Michael JD Powell. “A direct search optimization method that models the objective and constraint functions by linear interpolation”. In: *Advances in optimization and numerical analysis*. Springer, 1994, pp. 51–67.
- [270] Michael JD Powell. “The BOBYQA algorithm for bound constrained optimization without derivatives”. In: *Cambridge NA Report NA2009/06, University of Cambridge, Cambridge* (2009), pp. 26–46.
- [271] Michael JD Powell. “The NEWUOA software for unconstrained optimization without derivatives”. In: *Large-scale nonlinear optimization*. Springer, 2006, pp. 255–297.
- [272] Volker C Radeloff et al. “Rapid growth of the US wildland-urban interface raises wildfire risk”. In: *Proceedings of the National Academy of Sciences* 115.13 (2018), pp. 3314–3319.
- [273] Joaquin Ramirez, Santiago Monedero, and David Buckley. “New approaches in fire simulations analysis with wildfire analyst”. In: *7th International Conference on Forest Fire Research*. 2011.
- [274] Eghbal Rashidi, Hugh Medal, and Aaron Hoskins. “An attacker-defender model for analyzing the vulnerability of initial attack in wildfire suppression”. In: *Naval Research Logistics (NRL)* 65.2 (2018), pp. 120–134.
- [275] Eghbal Rashidi et al. “A maximal covering location-based model for analyzing the vulnerability of landscapes to wildfires: Assessing the worst-case scenario”. In: *European journal of operational research* 258.3 (2017), pp. 1095–1105.
- [276] Colleen E Reid et al. “Critical review of health impacts of wildfire smoke exposure”. In: *Environmental health perspectives* 124.9 (2016), pp. 1334–1343.
- [277] Elizabeth D Reinhardt and Matthew B Dickinson. “First-order fire effects models for land management: overview and issues”. In: *Fire Ecology* 6.1 (2010), pp. 131–142.
- [278] Fabienne Reisen and Stephen K Brown. “Implications for community health from exposure to bushfire air toxics”. In: *Environmental Chemistry* 3.4 (2006), pp. 235–243.
- [279] Francesco Ricci, Paolo Avesani, and Anna Perini. “Cases on fire: applying CBR to emergency management”. In: *NEW REV APPL EXPERT SYS* 5 (1999), pp. 175–190.
- [280] G. D. Richards. “An elliptical growth model of forest fire fronts and its numerical solution”. In: *Int. J. Numer. Math Eng* 30 (1990), pp. 1133–1149.
- [281] C Ricotta, D Guglietta, and A Migliozzi. “No evidence of increased fire risk due to agricultural land abandonment in Sardinia (Italy)”. In: *Natural Hazards and Earth System Sciences* 12.5 (2012), p. 1333.

- [282] Carlo Ricotta et al. “Assessing the influence of roads on fire ignition: Does land cover matter?” In: *Fire* 1.2 (2018), p. 24.
- [283] Marcos Rodrigues, Adrián Jiménez-Ruano, and Juan de la Riva. “Fire regime dynamics in mainland Spain. Part 1: drivers of change”. In: *Science of The Total Environment* (2019), p. 135841.
- [284] Brendan M Rogers et al. “Focus on changing fire regimes: interactions with climate, ecosystems, and society”. In: *Environmental Research Letters* 15.3 (2020), p. 030201.
- [285] Matthew G Rollins. “LANDFIRE: a nationally consistent vegetation, wildland fire, and fuel assessment”. In: *International Journal of Wildland Fire* 18.3 (2009), pp. 235–249.
- [286] R. C. Rothermel. “A mathematical model for predicting fire spread in wildland fires.” In: *US Dep. Agric., For. Serv., Intermt. For. Range Exp Stn Ogden, UT. Res. Pap. INT 115* (1972), p. 52.
- [287] Richard C Rothermel. “A mathematical model for predicting fire spread in wildland fuels”. In: *Res. Pap. INT-115. Ogden, UT: US Department of Agriculture, Intermountain Forest and Range Experiment Station. 40 p.* 115 (1972).
- [288] Francesco Rovero et al. “" Which camera trap type and how many do I need?" A review of camera features and study designs for a range of wildlife research applications.” In: *Hystrix* 24.2 (2013).
- [289] Francesco Rovero et al. “Estimating species richness and modelling habitat preferences of tropical forest mammals from camera trap data”. In: *PloS one* 9.7 (2014).
- [290] S. M. Running. “Is global warming causing more large wildfires?” In: *Science* 313 (2006), pp. 927–928. DOI: [10.1126/science](https://doi.org/10.1126/science).
- [291] Lucia Russo, Paola Russo, and Constantinos I Siettos. “A complex network theory approach for the spatial distribution of fire breaks in heterogeneous forest landscapes for the control of wildland fires”. In: *PloS one* 11.10 (2016), e0163226.
- [292] L Russo et al. “Complex network statistics to the design of fire breaks for the control of fire spreading”. In: *Chemical Engineering Transactions* (2015).
- [293] Adam Rytwinski and Kevin A Crowe. “A simulation-optimization model for selecting the location of fuel-breaks to minimize expected losses from forest fires”. In: *Forest ecology and management* 260.1 (2010), pp. 1–11.
- [294] Muñoz Sabater and J ERA5-Land Monthly Averaged Data. “from 1981 to Present”. In: *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)* (2019).
- [295] Stavros Sakellariou et al. “Spatial planning of fire-agency stations as a function of wildfire likelihood in Thasos, Greece”. In: *Science of the total environment* 729 (2020), p. 139004.

- [296] Michele Salis et al. “Evaluating alternative fuel treatment strategies to reduce wildfire losses in a Mediterranean area”. In: *Forest Ecology and Management* 368 (2016), pp. 207–221.
- [297] Michele Salis et al. “Modeling the effects of different fuel treatment mosaics on wildfire spread and behavior in a Mediterranean agro-pastoral area”. In: *Journal of environmental management* 212 (2018), pp. 490–505.
- [298] J. C. Sanderlin and J. M. Sunderson. “A simulation for wildland fire management planning support (FIREMAN). Vol. 2”. In: *Prototype models for FIREMAN (Part II): Campaign fire evaluation Mission Research Corp. Contract 21-343, Spec. 222 2* (1975), p. 249.
- [299] Pablo Sarricolea et al. “Recent wildfires in Central Chile: Detecting links between burned areas and population exposure in the wildland urban interface”. In: *Science of the Total Environment* 706 (2020), p. 135894.
- [300] Boris Schäling. *The boost C++ libraries*. Boris Schäling, 2011.
- [301] Erich Schubert et al. “DBSCAN revisited, revisited: why and how you should (still) use DBSCAN”. In: *ACM Transactions on Database Systems (TODS)* 42.3 (2017), pp. 1–21.
- [302] Ronald L Schumann III et al. “Wildfire recovery as a “hot moment” for creating fire-adapted communities”. In: *International journal of disaster risk reduction* 42 (2020), p. 101354.
- [303] Dylan W Schwilk and Benjamin Kerr. “Genetic niche-hiking: an alternative explanation for the evolution of flammability”. In: *Oikos* 99.3 (2002), pp. 431–442.
- [304] Andrew C Scott and Ian J Glasspool. “The diversification of Paleozoic fire systems and fluctuations in atmospheric oxygen concentration”. In: *Proceedings of the National Academy of Sciences* 103.29 (2006), pp. 10861–10865.
- [305] Andrew C Scott et al. *Fire on earth: an introduction*. John Wiley & Sons, 2013.
- [306] J. H. Scott and R. E. Burgan. “Standard fire behavior fuel models: a comprehensive set for use with Rothermel’s surface fire spread model.” In: *US Dep. Agric., For. Serv., Rocky Mtn. Res. Stn., Fort Collins, CO Tech. Rep. RMRS-GTR-153* (2005), p. 72.
- [307] JH Scott and RE Burgan. *Standard fire behavior fuel models: a comprehensive set for use with Rothermel’s surface fire spread model. USDA Forest Service, Rocky Mountain Research Station*. Tech. rep. General Technical Report RMRS-GTR-153.(Fort Collins, CO), 2005.
- [308] Terrence J Sejnowski. “The unreasonable effectiveness of deep learning in artificial intelligence”. In: *Proceedings of the National Academy of Sciences* (2020).

- [309] Ramprasaath R Selvaraju et al. “Grad-cam: Visual explanations from deep networks via gradient-based localization”. In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 618–626.
- [310] John Sessions et al. “Achieving sustainable forest structures on fire-prone landscapes while pursuing multiple goals”. In: *Spatial Modeling of Forest Landscape Change: Approaches and Applications*. Cambridge University Press, Cambridge, UK (1999), pp. 210–253.
- [311] Antoine Sevilla and Hervé Glotin. “Audio Bird Classification with Inception-v4 extended with Time and Time-Frequency Attention Mechanisms.” In: *CLEF (Working Notes)*. 2017.
- [312] Douglas J Shinneman et al. “The ecological uncertainty of wildfire fuel breaks: examples from the sagebrush steppe”. In: *Frontiers in Ecology and the Environment* (2019).
- [313] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. “Learning important features through propagating activation differences”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 3145–3153.
- [314] Karen Simonyan and Andrew Zisserman. “Very deep convolutional networks for large-scale image recognition”. In: *arXiv preprint arXiv:1409.1556* (2014).
- [315] Amber J Soja et al. “Climate-induced boreal forest change: predictions versus current observations”. In: *Global and Planetary Change* 56.3-4 (2007), pp. 274–296.
- [316] Maria João Sousa, Alexandra Moutinho, and Miguel Almeida. “Wildfire detection using transfer learning on augmented datasets”. In: *Expert Systems with Applications* 142 (2020), p. 112975.
- [317] WGRFR Spendley, George R Hext, and Francis R Himsworth. “Sequential application of simplex designs in optimisation and evolutionary operation”. In: *Technometrics* 4.4 (1962), pp. 441–461.
- [318] Zachary L Steel, Michael J Koontz, and Hugh D Safford. “The changing landscape of wildfire: burn pattern trends and implications for California’s yellow pine and mixed conifer forests”. In: *Landscape Ecology* 33.7 (2018), pp. 1159–1176.
- [319] Chris Stockdale et al. “Examining management scenarios to mitigate wildfire hazard to caribou conservation projects using burn probability modeling”. In: *Journal of environmental management* 233 (2019), pp. 238–248.
- [320] B. J. Stocks and D. L. Martell. “Forest fire management expenditures in Canada: 1970–2013”. In: *The Forestry Chronicle* 92.3 (2016), pp. 298–306.
- [321] Brian R Sturtevant et al. “Simulating dynamic and mixed-severity fire regimes: a process-based fire extension for LANDIS-II”. In: *Ecological Modelling* 220.23 (2009), pp. 3380–3393.

- [322] Andrew L Sullivan. “Wildland surface fire spread modelling, 1990–2007. 1: Physical and quasi-physical models”. In: *International Journal of Wildland Fire* 18.4 (2009), pp. 349–368.
- [323] K Sun et al. “Wildfire spread in Wildland Urban Interface (WUI) communities in California: Introducing the urban fuel”. In: *AGUFM 2019* (2019), NH43C–0950.
- [324] Richard S Sutton and Andrew G Barto. “Reinforcement learning: An introduction”. In: (2011).
- [325] Alexandra D Syphard and Jon E Keeley. “Location, timing and extent of wildfire vary by cause of ignition”. In: *International Journal of Wildland Fire* 24.1 (2015), pp. 37–47.
- [326] Alexandra D Syphard et al. “Human influence on California fire regimes”. In: *Ecological applications* 17.5 (2007), pp. 1388–1402.
- [327] Ming Tan. “Multi-agent reinforcement learning: Independent vs. cooperative agents”. In: *Proceedings of the tenth international conference on machine learning*. 1993, pp. 330–337.
- [328] Shawn Tan, Khe Chai Sim, and Mark Gales. “Improving the interpretability of deep neural networks with stimulated learning”. In: *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE. 2015, pp. 617–623.
- [329] Stephen W Taylor and Martin E Alexander. “Science, technology, and human factors in fire danger rating: the Canadian experience.” In: *International Journal of Wildland Fire* 15.1 (2006), pp. 121–135.
- [330] Tippaya Thinsungnoena et al. “The clustering validity with silhouette and sum of squared errors”. In: *learning* 3.7 (2015).
- [331] Matthew P Thompson et al. “A review of challenges to determining and demonstrating efficiency of large fire management”. In: *International Journal of Wildland Fire* 26.7 (2017), pp. 562–573.
- [332] Matthew P Thompson et al. “Integrated wildfire risk assessment: framework development and application on the Lewis and Clark National Forest in Montana, USA”. In: *Integrated Environmental Assessment and Management* 9.2 (2013), pp. 329–342.
- [333] Xinmin Tian et al. “Intel® OpenMP C++/Fortran Compiler for Hyper-Threading Technology: Implementation and Performance.” In: *Intel Technology Journal* 6.1 (2002).
- [334] B. Todd. “User documentation for the Wildland Fire Growth Model and the Wildfire Display Program. Can. For. Serv., North For”. In: *Cent., Edmonton, AB* 37 p. Draft rep (1999).
- [335] Juan J Troncoso, Andrés Weintraub, and David L Martell. “Development of a threat index to manage timber production on flammable forest landscapes subject to spatial harvest constraints”. In: *INFOR: Information Systems and Operational Research* 54.3 (2016), pp. 262–281.

- [336] Merritt R Turetsky et al. “Global vulnerability of peatlands to fire and carbon loss”. In: *Nature Geoscience* 8.1 (2015), p. 11.
- [337] Monica G Turner and Robert H Gardner. *Quantitative methods in landscape ecology: the analysis and interpretation of landscape heterogeneity*. 574.5 T8. 1991.
- [338] C. Tymstra et al. “Development and structure of Prometheus: the Canadian Wildland Fire Growth Simulation Model”. In: *Information Report NOR-X-Edmonton (AB): Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre* 417 (2010), p. 102.
- [339] C Tymstra et al. “Development and structure of Prometheus: the Canadian wildland fire growth simulation model”. In: *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Information Report NOR-X-417.(Edmonton, AB)* (2010).
- [340] Cordy Tymstra et al. “Wildfire management in Canada: Review, challenges and opportunities”. In: *Progress in Disaster Science* 5 (2020), p. 100045.
- [341] Alfred Ultsch and H. Peter Siemon. “Kohonen’s Self Organizing Feature Maps for Exploratory Data Analysis”. In: *Proceedings of the International Neural Network Conference (INNC-90), Paris, France, July 9–13, 1990 1*. Dordrecht, Netherlands. Ed. by Bernard Widrow and Bernard Angeniol. Vol. 1. Dordrecht, Netherlands: Kluwer Academic Press, 1990, pp. 305–308. URL: <http://www.uni-marburg.de/fb12/datenbionik/pdf/pubs/1990/UltschSiemon90>.
- [342] Evelyn Uuemaa, Ülo Mander, and Riho Marja. “Trends in the use of landscape spatial metrics as landscape indicators: a review”. In: *Ecological Indicators* 28 (2013), pp. 100–106.
- [343] Guido R Van der Werf et al. “Global fire emissions and the contribution of deforestation, savanna, forest, agricultural, and peat fires (1997-2009)”. In: *Atmospheric Chemistry and Physics* 10.23 (2010), pp. 11707–11735.
- [344] Hado Van Hasselt, Arthur Guez, and David Silver. “Deep reinforcement learning with double q-learning”. In: *Thirtieth AAAI conference on artificial intelligence*. 2016.
- [345] C. E. Van Wagner and T. L. Pickett. *Development and structure of the Canadian Forest Fire Weather Index System*. Technical Report 35. Canadian Forest Service, Forestry (Ottawa, ON), 1987.
- [346] Jan W Van Wagendonk. “Fire as a physical process”. In: *Fire in California’s ecosystems*. University of California Press, Berkeley, USA (2006), pp. 38–57.
- [347] Boris Vanni ere et al. “Climate versus human-driven fire regimes in Mediterranean landscapes: the Holocene record of Lago dell’Accesa (Tuscany, Italy)”. In: *Quaternary Science Reviews* 27.11-12 (2008), pp. 1181–1196.

- [348] Olga Viedma, Jose M Moreno, and Ignacio Rieiro. “Interactions between land use/land cover change, forest fires and landscape structure in Sierra de Gredos (Central Spain)”. In: *Environmental Conservation* (2006), pp. 212–222.
- [349] Kevin C Vogler et al. “Prioritization of forest restoration projects: tradeoffs between wildfire protection, ecological restoration and economic objectives”. In: *Forests* 6.12 (2015), pp. 4403–4420.
- [350] Dale R Waddell. “Estimating load weights with Huber’s Cubic Volume formula: a field trial.” In: *Res. Note. PNW-RN-484. Portland, OR: US Department of Agriculture, Forest Service, Pacific Northwest Research Station. 12 p* 484 (1989).
- [351] CE Van Wagner. “A simple fire-growth model”. In: *The Forestry Chronicle* 45.2 (1969), pp. 103–104.
- [352] CE Van Wagner. “Conditions for the start and spread of crown fire”. In: *Canadian Journal of Forest Research* 7.1 (1977), pp. 23–34.
- [353] CE Van Wagner. “Height of crown scorch in forest fires”. In: *Canadian journal of forest research* 3.3 (1973), pp. 373–378.
- [354] Zhou Wang et al. “Image quality assessment: from error visibility to structural similarity”. In: *IEEE transactions on image processing* 13.4 (2004), pp. 600–612.
- [355] Christopher JCH Watkins and Peter Dayan. “Q-learning”. In: *Machine learning* 8.3-4 (1992), pp. 279–292.
- [356] Yu Wei. “Optimize landscape fuel treatment locations to create control opportunities for future fires”. In: *Canadian Journal of Forest Research* 42.6 (2012), pp. 1002–1014.
- [357] Yu Wei, Douglas B Rideout, and Thomas B Hall. “Toward efficient management of large fires: a mixed integer programming model and two iterative approaches”. In: *Forest Science* 57.5 (2011), pp. 435–447.
- [358] Yu Wei, Douglas Rideout, and Andy Kirsch. “An optimization model for locating fuel treatments across a landscape to reduce expected fire losses”. In: *Canadian Journal of Forest Research* 38.4 (2008), pp. 868–877.
- [359] A. L. Westerling. “Increasing western US forest wildfire activity: sensitivity to changes in the timing of spring”. In: *Phil. Trans. R. Soc. B* 371.20150178 (2016).
- [360] A. L. Westerling et al. “Warming and earlier spring increase western US forest wildfire activity”. In: *Science* 313 (2006), pp. 940–943. DOI: [10.1126/science.1128834](https://doi.org/10.1126/science.1128834).
- [361] Cathy Whitlock et al. “Paleoecological perspectives on fire ecology: revisiting the fire-regime concept”. In: *The Open Ecology Journal* 3.1 (2010).
- [362] Marco A Wiering and Marco Dorigo. “Learning to control forest fires”. In: (1998).
- [363] A Park Williams and John T Abatzoglou. “Recent advances and remaining uncertainties in resolving past and future climate effects on global fire activity”. In: *Current Climate Change Reports* 2.1 (2016), pp. 1–14.

- [364] Svante Wold, Kim Esbensen, and Paul Geladi. “Principal component analysis”. In: *Chemometrics and intelligent laboratory systems* 2.1-3 (1987), pp. 37–52.
- [365] David H Wolpert and Kagan Tumer. “Optimal payoff functions for members of collectives”. In: *Modeling complexity in economic and social systems*. World Scientific, 2002, pp. 355–369.
- [366] Dayong Ye, Minjie Zhang, and Yun Yang. “A multi-agent framework for packet routing in wireless sensor networks”. In: *sensors* 15.5 (2015), pp. 10026–10047.
- [367] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. “Recurrent neural network regularization”. In: *arXiv preprint arXiv:1409.2329* (2014).
- [368] Guoli Zhang, Ming Wang, and Kai Liu. “Forest Fire Susceptibility Modeling Using a Convolutional Neural Network for Yunnan Province of China”. In: *International Journal of Disaster Risk Science* 10.3 (2019), pp. 386–403.
- [369] Pengfei Zhang, Huitao Shen, and Hui Zhai. “Machine learning topological invariants with neural networks”. In: *Physical review letters* 120.6 (2018), p. 066401.
- [370] Quan-shi Zhang and Song-Chun Zhu. “Visual interpretability for deep learning: a survey”. In: *Frontiers of Information Technology & Electronic Engineering* 19.1 (2018), pp. 27–39.
- [371] Shizhou Zhang et al. “Constructing deep sparse coding network for image classification”. In: *Pattern Recognition* 64 (2017), pp. 130–140.
- [372] Yuanyuan Zhao et al. “Detailed dynamic land cover mapping of Chile: Accuracy improvement by integrating multi-temporal data”. In: *Remote Sensing of Environment* 183 (2016), pp. 170–185.
- [373] W. Zhou et al. “Image Quality Assessment: From Error Visibility to Structural Similarity”. In: *IEEE Transactions on Image Processing* 13.4 (Apr. 2004), pp. 600–612.

Appendix A

Cell2Fire: Cellular-Automata Wildfire Simulator

A.1 Cells attributes

Cells have attributes concerning the structure of the forest as well as the past and current state of the simulation are the main processing units in Cell2Fire. The main fields that define the state of a cell $i \in \mathcal{N}$ are the following:

- i) Fuel type: following the classification criteria of the Canadian FBP System, each cell is assigned a specific fuel type (e.g., conifer, grass, non-fuel) represented by an FBP code, which will be used for selecting the specific fire spread models and corresponding coefficients defined by the FBP System in order to predict the ROS in that cell.
- ii) Slope: the slope (%) and the upslope direction (radians) have a significant impact on the predicted ROS.
- iii) Elevation: altitude in meters of the current cell with respect to the sea level.
- iv) Location: latitude and longitude coordinates are provided with the instance.
- v) Status: cells are classified into five different states during the simulation time steps, namely “*Available*”, “*Burning*”, “*Burned*”, “*Treated*”, or “*Non-Fuel*”.

In addition, each cell contains a series of secondary parameters that allow Cell2Fire to track the growth of the fire within the forest and change the state of the current simulation run. The fire dynamics group includes the fire ignition date and time of each cell, the fire’s progress along each axis, and the effective rate of spread $ROS(t, \theta)$ values (per period and axis angle). Due to this independent structure, cells can be treated as individual units allowing an efficient parallel computing approach for each time step. We can therefore update their status and generate the relevant fire messages to model the fire dynamics of each burning cell in parallel using independent threads and thereby obtain significant improvements in

execution times (from hours in serial mode to minutes in parallel mode) when dealing with large fires that have many simultaneously active cells.

A.2 Fire propagation dynamics

The fire growth model is straightforward. Every time a cell is ignited by an adjacent cell it acts as a new source of potential ignition for neighboring cells in the forest, updating the progress of the fire along each available axis (center-to-center directions). Once the fire instance data has been read and the forest has been initialized, the main simulation steps are as follows:

Algorithm 2 Cell2Fire Pseudo-code

```

1: procedure SIM(ForestData, FTypes, Ignitions, Weather, TMax, Options)
2:   Step 0: Initialize Cell2Fire
3:     Parse inputs, options, read data, initialize objects
4:      $FPeriod \leftarrow 0, n_{sim} \leftarrow 0$ 
5:   Step 1: Ignition
6:      $ic \leftarrow \text{Choice}(\text{Ignitions})$ 
7:      $BurningCells \cup \{ic\}$ 
8:      $FPeriod \leftarrow 1$ 
9:   Step 2: Fire Dynamics (Send-Receive)
10:    Let  $i \in BurningCells$ , if  $i \rightarrow_{FPeriod} j$ , where  $j \in Adj(i)$ :
11:       $BurningCells \cup \{j\}$ 
12:     $FPeriod \leftarrow FPeriod + 1$ , Update Forest, Weather
13:    Repeat, until  $FPeriod > TMax$ 
14:  Step 3: Results and Outputs generation
15:    Generate Grids, Statistics, Output plots
16: end procedure=0

```

where we use the following notation to describe the main steps of the simulation shown in Algorithm 2:

- \mathcal{N} denotes the set of cells in the forest and $i \in \mathcal{N}$ represents cell $i \in \mathcal{N}$.
- $Adj(i)$ denotes the set of available cells that are adjacent to cell i (at most 8).
- $BurningCells$ is the set of cells that are currently burning.
- $i \rightarrow_t j$ indicates that cell i “sends a message” to cell j at time t

i) Relevant fire parameters are calculated by performing calls to the FBP System module to determine the ROS for each available fire spread axis of burning cells based on

fuel characteristics, topography, and weather. Following a discrete time simulation approach, the internal simulator clock advances one unit of time – a user-input precision parameter – and the fire’s progress is updated along each axis for each burning cell.

- ii) Fire spread between cells is modeled by using a sending/receiving message approach (which enables parallelization) based on the computed ROS along each axis. If the fire reaches the center of an adjacent cell during the simulated time step, a message is sent and instantaneously received. Checking environmental conditions and the cell’s characteristics, the cell begins burning (or not). This process is the core of the simulator and the critical performance bottleneck that comes into play when simulating fire spread across large landscapes. However, we designed it to maximize the parallel performance of the code.
- iii) The previous steps are repeated until some specified ending criterion is satisfied: e.g., the maximum number of weather periods, the maximum simulation time, and/or some fire-ending event condition. Statistics regarding the status of the forest as well as plots and other outputs of the fire scar evolution are produced.

Rate of Spread from cell to cell

We use the Canadian FBP System and the procedure for parameterizing an elliptical fire growth model described to model the growth of elliptical fires in each burning cell. This method does not require the ignition point or the point of propagation to coincide with either of the two foci of the ellipse, although the authors indicate that “*small differences between the point of ignition and the focus of the ellipse do not change the results*”, i.e., the elliptical propagation shape/size. The components of the fire spread model are the following:

- An elliptical fire has a HROS (head), a BROS (backing) and a FROS (flank). Let a , b , and e be the semi-major, semi-minor, and eccentricity of the ellipse, respectively.
- The FBP system predicts the HROS, the BROS and the length to breadth ratio (LB) which is $2a/2b$ or a/b .
- During the first time interval t , the fire will spread from its ignition point towards the center of the ellipse and then from the center of the ellipse to the farthest edge of it.
- At time t , we have:

$$a = \frac{HROS + BROS}{2} \times t, \quad (\text{A.1})$$

$$b = \frac{2 \times FROS}{2} \times t, \quad (\text{A.2})$$

to expand the ellipse generated by the propagation of fire at time t on its two main axes.

- Noting that $LB = \frac{a}{b}$ we have $FROS = \frac{HROS + BROS}{2LB}$. Therefore, the eccentricity is:

$$e = \sqrt{1 - \left(\frac{(FROS \times t)^2}{\frac{(HROS + BROS) \times t^2}{2}} \right)} \quad (\text{A.3})$$

Using these equations and the procedure described in Prometheus, we can estimate the ROS from the center of a cell to the center of any adjacent cell as in equation (1).

A.3 Detailed inputs

Scenario	datetime	APCP [mm]	TMP [C°]	RH [%]	WS [m/s]	WD[°]
JCB	2001-10-16 13:00	0.0	17.7	20	21	225
JCB	2001-10-16 14:00	0.6	16.9	18	25	205
JCB	2001-10-16 15:00	1.2	16.1	20	27	190
JCB	2001-10-16 16:00	10.0	15.8	20	37	232
JCB	2001-10-16 17:00	5.3	13.9	25	43	225
JCB	2001-10-16 18:00	2.1	12.1	35	45	222
JCB	2001-10-16 19:00	0.9	10.6	41	46	241
JCB	2001-10-16 20:00	0.0	11.3	39	18	248

Table A.1: Extract of an hourly weather stream file. Average precipitation (APCP), temperature (TMP), relative humidity (RH), wind speed (WS) and wind direction (WD).

Sample files are included with the distribution of Cell2Fire for the publicly available Dogrib fire instance¹ as well as simple generated test instances. In addition to the main input files, a set of options and user-provided parameters for exploiting all the flexibility of the simulation engine including adjusting options for the ROS are available as secondary inputs when running Cell2Fire.

¹http://www.firegrowthmodel.ca/prometheus/software_e.php

A.4 Propagation model validation

Dogrib fire

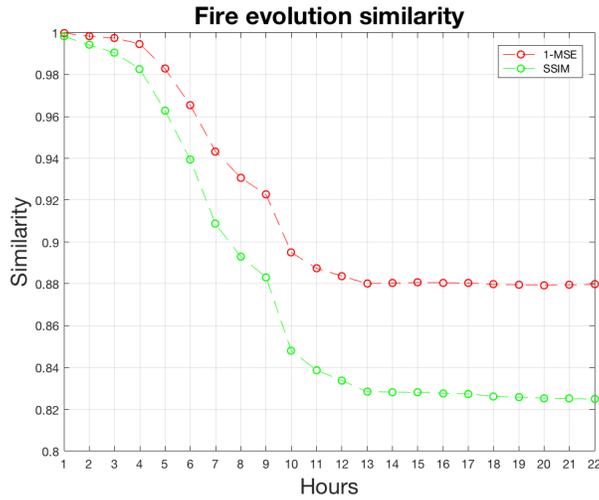


Figure A.1: Dogrib MSE and SSIM hourly evolution (22 hours simulation).

Hour	1-MSE [%]	SSIM [%]	Hour	1-MSE [%]	SSIM [%]
1	99.98	99.83	12	88.37	83.38
2	99.85	99.42	13	88.01	82.84
3	99.74	99.05	14	88.03	82.83
4	99.47	98.28	15	88.06	82.81
5	98.29	96.27	16	88.06	82.77
6	96.54	93.94	17	88.04	82.74
7	94.33	90.86	18	87.98	82.62
8	93.08	89.30	19	87.96	82.58
9	92.26	88.32	20	87.94	82.54
10	89.50	84.81	21	87.95	82.51
11	88.74	83.87	22	87.97	82.49
AVG [%]	95.62	93.08	AVG [%]	88.03	82.74

Table A.2: Dogrib accuracy and structural similarity index measure values per hour (22 hours evolution).

A.5 Parallel performance analysis

Methodology

The high complexity of the simulation scheme poses a challenge for measuring the performance of our implementation: each fire has a large number of parameters and specific characteristics that lead to different outcomes. Thus, different fuel types, weather streams, forest structures, and/or ignition points could lead to completely different fire dynamics (e.g., number of simultaneous burning cells, number of burned hectares), and hence, to different performance when operating serially or in parallel. In order to account for this, multiple instances/forests based on real fuel, weather and topographic data are generated and average results are used to describe the performance of Cell2Fire.

Performance is measured by calculating both the strong and weak scaling efficiencies – as well as speedup factors – obtained for different experimental instances ranging from sizes (number of cells inside the forest) $n \in [4 - 1M]$. Summary plots are generated in order to visualize the performance of our parallel implementation.

Instances: Data

Two weather files from weather stations located in Canada containing all relevant inputs for 7 and 36 hours were used for all instances. Each set of experiments for n cells uses the same ignition points for comparison purposes, starting fire growth at the same time for 1-minute time-steps. Larger instances are generated by mixing locations using real landscape attribute data gathered from Canadian forests. In addition, homogeneous instances (same fuel type for all cells) are included in each experimental set for the purpose of performance experiments.

Hardware & Software

The optimization and parallelization of Cell2Fire were developed for a specific hardware and run-time environment for the National Energy Research Scientific Computing Center (NERSC). In addition, tests have been performed in a common laptop for comparison purposes. All experiments, benchmarks, and performance results are implemented using the following hardware and software:

1. NERSC's Cori supercomputer: Phase I

- Intel® Xeon™ Processor E5-2698 v3 ("Haswell") at 2.3 GHz (32 cores per node)
- 64 KB 8-way L1, 256 KB 8-way set L2, and 40 MB 20-way set L3 cache (shared per socket)
- SUSE Linux version 4.4.74-92.38-default. Built with g++ version 4.8.5

2. Laptop

- Intel® Core I7 4510U at 2.0 GHz (2 cores)
- 64 KB 8-way set L1, 2 x 256 KB 8-way set L2, and 4 MB 16-way set L3 cache
- Ubuntu 16.04.2 LTS / Windows 10

Parallel Structure

Our algorithm performs three operations at each time step: (1) checking for new lightning ignitions (igniting), (2) updating the fire progress of already-burned cells and analyzing newly burned ones (sending messages), and (3) marking newly burned cells as burning (receiving messages).

The ignition stage is very quick (less than 1% of the total execution time), with most simulations only igniting a single time at the first time step of the simulation. The sending messages stage updates the ellipse associated with every burning cell. Because we can have a large number of cells burning at once and there are no direct dependencies on neighboring cells, this part is easily parallelizable. Each cell, in addition to updating itself, can also “send a burning message” to an adjacent cell. In the receiving messages stage, we analyze the “burn messages” sent to non-burning cells and mark them as burned if the conditions are met. This part is also potentially parallelizable, but because the number of newly burned cells at a single time-step is dwarfed by the number of currently burning cells, we found that a speedup here is of lower priority ($\approx 10\%$ of total execution time).

Parallelization: OpenMP work-sharing

Due to the easily parallelizable structure of our code, the most suitable approach for parallelizing its execution consists of a shared-memory approach using the well-known OpenMP API. This is an advantage since the code will be also optimized for its execution in normal desktop/laptop computers without needing a multi-node architecture to exploit parallelism. We found that we could easily make the loop embarrassingly parallel if instead of adding to a single data structure, we add to a data structure local to the worker thread. Since we would iterate over the initial data structure to compute statistics, we found the additional complexity of “distributing” the data structure does not scale with grid size. In addition, different loop scheduling options were tested: dynamic, guided, auto, run-time, and static, as well as the chunk-sized block process. Following a brute force optimization approach, we were able to obtain an average of 15%-20% extra performance for the parallel region.

One final improvement we made to our parallelism was analyzing the false-sharing effect. Because we had a *vector* $\langle DS \rangle$ to store our “distributed” data structures, where DS is the data structure of choice, we found there to be a bottleneck on the parallelism exploited in the problem. Upon further analysis, we found this to be false-sharing in the array of DS backing the vector. After adding padding between elements of the array – where optimal values

were obtained following a binary search optimization approach – we achieved a significant additional speedup from our initial attempt.

Parallel section: percentage of the total code

In order to identify the potential benefits of a parallelization, we performed a detailed analysis of the execution times, breaking it into: (1) sending time (“parallel” region), (2) receiving time, (3) ignition time, and (4) copying time. In Tables A.3 (a), (b) we can see that the average — across all instances — time spent in the parallelizable region represents a 79% of the total execution time. The other $\approx 20\%$ is divided evenly between (2) and (4), while (3) is almost negligible. These results give us a sense of the potential impact of an efficient parallel implementation in our code.

Instance (n)	AVG % Time in Parallel Zone: Sending			Instance (n)	AVG % Time in Parallel Zone: Sending		
	Bottom 10%	Middle 80%	Top 10%		Bottom 10%	Middle 80%	Top 10%
4	80.44%	86.00%	94.10%	160K	72.43%	87.60%	88.32%
9	74.15%	82.63%	93.11%	250K	70.82%	78.28%	90.05%
400	56.75%	72.09%	95.84%	500K	69.94%	73.70%	80.19%
1600	64.32%	71.57%	90.45%	1M	68.64%	78.01%	89.00%
10K	54.00%	77.46%	88.34%	10M	60.25%	80.00%	92.43%
50K	62.00%	71.73%	87.43%	50M	58.30%	88.77%	93.11%
100K	64.50%	79.13%	91.23%	100M	52.00%	90.03%	95.24%
AVG	65.17%	77.23%	91.50%	AVG	64.63%	82.34%	89.76%

Table A.3: Average percentage of the running time parallelizable for different instances. Upper and lower tails are included for completeness. Results obtained by simulating 20 different forests (fuel types, spatial distribution, etc.) for each size n with the same weather conditions.

It is important is to note the fact that some instances (lower tail) will experience a poor parallel performance. The reason behind this behavior is clear: certain combinations of fuel types and forests distributions lead to a significantly smaller set of simultaneous burning cells per simulation, and thus, the parallelization of (1) will not impact the overall execution times as much as we desire. Knowing this limitation, we proceed to analyze the solving times and accuracy of the simulations.

Running times & speedup analysis

Comparing the running times of our pure Python prototype and C++ implementations with the results obtained using Prometheus, we can see in Figure A.2, how the optimized version clearly outperforms the pure Python prototype, reaching up to 15-20x speedups when dealing with large instances. Furthermore, Cell2Fire (C++) obtains significant shorter times than Prometheus (up to 30x speedups). This is important since wave-based simulators perform a

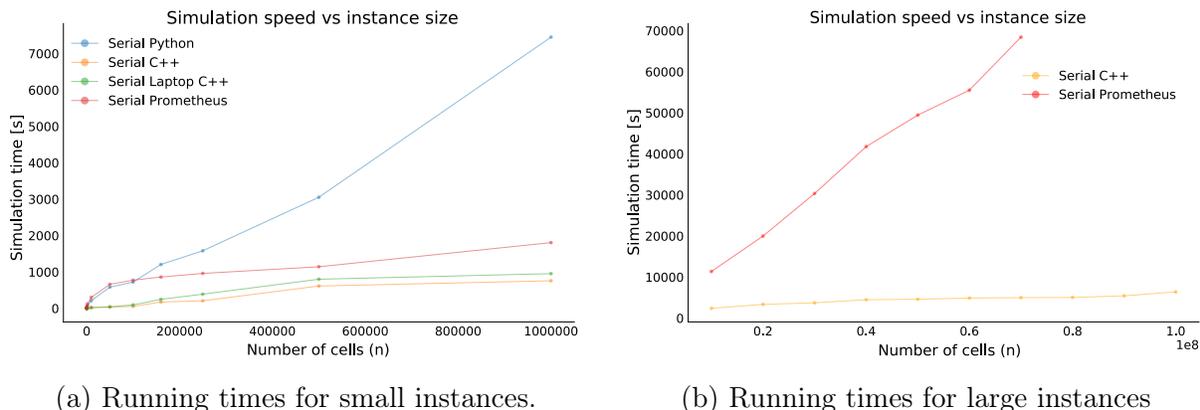


Figure A.2: Running times for serial versions. Python’s large instances results are omitted for visualization purposes.

series of approximations when generating the final fire scar that simplifies the calculations performed under our cell-based approach, indicating that our implementation is efficient.

It is interesting to note that Prometheus is not able to solve our three largest instances (80M, 90M, and 100M, due to an "out of memory" error). Looking at Table A.4 we can see that the detailed and average speedups obtained for the small instances with the optimized OpenMP version are very good, obtaining high-performance with a certain number of threads. Performance is even better when dealing with the large instances, improving each average speedup up to an average of 16.48x when running 32 parallel threads.

Table A.4: Speedup factors for small instances for different numbers of threads.

Instance (n)	2 threads	4 threads	8 threads	16 threads	24 threads	32 threads
4	0.50	0.40	1.00	1.00	1.50	6.00
9	0.71	2.50	5.00	3.33	5.00	5.00
400	3.01	3.26	3.69	3.66	3.66	3.69
1600	3.07	3.50	4.18	4.34	4.34	4.34
10000	2.24	2.81	3.64	4.84	6.12	8.70
50000	1.93	3.21	3.51	3.86	4.23	8.66
100000	1.83	3.72	4.26	4.35	5.71	8.38
160000	1.86	3.05	3.29	3.63	3.58	10.50
250000	1.85	3.21	4.76	5.52	9.12	11.98
500000	1.98	3.71	6.44	8.88	13.06	17.45
1000000	2.10	3.67	6.09	8.99	12.16	17.13
AVG Speedup	1.92	3.00	4.17	4.76	6.23	9.26

As expected, even better speedups are obtained when the dealing with homogeneous forests as can be seen in the Table A.5 where a summary for the large instances speedup averages is shown. From this, we can see a near optimal average speedup up to 16 threads, while reaching a great $\approx 20x$ with 32 threads.

Table A.5: Average speedups for large instances: heterogeneous and homogeneous forests

	2 threads	4 threads	8 threads	16 threads	24 threads	32 threads
AVG Speedup Large Homogeneous	1.99	3.76	7.01	12.33	15.34	19.78
AVG Speedup Large Heterogeneous	1.84	2.44	4.22	6.89	11.62	16.48

Strong Scaling

After generating the speedup and strong-scaling efficiency plots for the experimental instances, we observe that the optimized implementation is able to obtain up to 15x and 20x speedups for the small and large instances respectively, as well as averages strong efficiency factors between 75% and 82%, depending on the size and structure of the forest. In Figures A.3 (a) and (b) we present the results obtained for the average values obtained among 20 instances with 500,000 cells using the optimized OpenMP implementation. Similar — slightly better — results are obtained for larger instances.

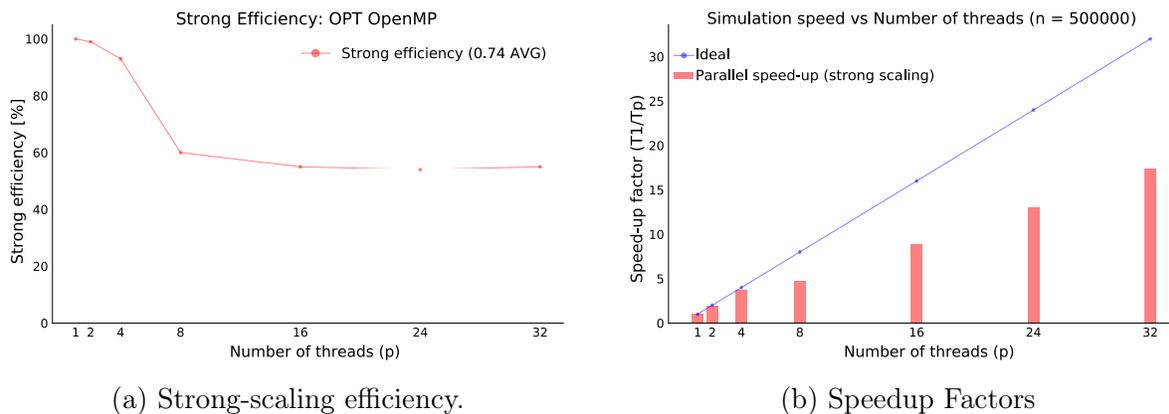


Figure A.3: Strong-scaling and speedup factors for OPT version $n = 500,000$.

Based on all our experiments, adding more threads leads to better execution times following a flat pattern w.r.t. the strong-scaling efficiency. Thus, our optimized implementation is able to obtain a great strong scaling performance, taking into account the high complexity of the instances and variability of the results depending on the forest's structure.

Weak Scaling

Due to the high dependency of the results on forest structure, we separated our analysis for heterogeneous and homogeneous forests. When the forests are heterogeneous, the comparison between instances of different sizes and number of threads loses its meaning since there is no guarantee that the problem will scale in complexity: increasing the problem size does not affect the computations time directly, it depends significantly on the composition of the forest, leading to different fire dynamics. Therefore, we expect an erratic pattern when dealing with heterogeneous forests in terms of weak-scaling efficiency.

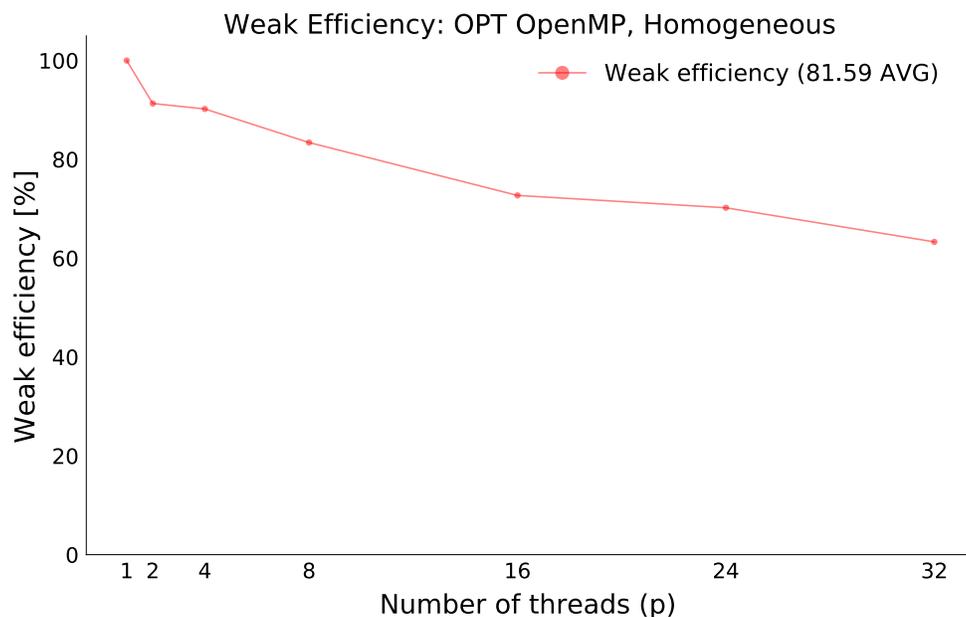


Figure A.4: Weak-scaling efficiency (Homogeneous)

Figure A.5: Weak-scaling and speedup factors for homogeneous instances, starting with $n = 500,000$.

In Figure A.5 we can see the weak efficiency obtained for homogeneous instances. The results for heterogeneous instances lack meaning since we are not comparing the same fire dynamics (and thus, the number of simultaneous burning cells, critical for the parallel performance). On the other hand, comparing homogeneous instances gives us correct results, since we compare the same fire dynamics reaching an average weak-scaling factor value equal to 81.6%. Again, similar and even better weak factors are obtained with larger (and homogeneous) instances, following the discussion above.

Therefore, our optimized parallel implementation is able to obtain high-performance values in both strong and weak scaling factors thanks to its naturally parallel design.

Appendix B

General Wildfire Spread Model

B.1 Additional fire components

Interaction with neighboring cells

The ROS $r(\phi)$ function can be calculated for all angles ϕ using:

$$r(\phi) = a(1 - e^2) / (1 - e \cdot \cos \phi) \quad (\text{B.1})$$

Cell2Fire calculates the ROS towards the 8 directions where the neighboring cells are located. We denote by ROS_{ij} the fire that travels from cell i to cell j and let $Adj(i)$ the set of adjacent cells to i . Each adjacent cell to i corresponds to an angle $\phi \in \{0^\circ, \dots, 315^\circ\}$ by a step of 45° . Thus, $ROS_{i,i_5} = r(0^\circ)$, $ROS_{i,i_3} = r(45^\circ)$, $ROS_{i,i_2} = r(90^\circ)$, $ROS_{i,i_1} = r(135^\circ)$, $ROS_{i,i_4} = r(180^\circ)$, $ROS_{i,i_6} = r(225^\circ)$, $ROS_{i,i_7} = r(270^\circ)$, and $ROS_{i,i_8} = r(315^\circ)$.

In the propagation dynamics of Cell2Fire, once the ROS_{ij} values have been calculated, the effect of the slope (SE) is introduced by means of a simple equation extracted from the Kitral system [170].

$$SE_{ij} = 1.0 + 0.023322 \cdot SLOPE_{ij} + 0.00013585 \cdot SLOPE_{ij}^2 \quad (\text{B.2})$$

where the $SLOPE_{ij}$ value is the slope between adjacent cells i and j and their corresponding elevations $ELEV_i$ and $ELEV_j$, where $SLOPE_{ij} = (ELEV_j - ELEV_i)/L$, with L the length of an individual cell. Therefore, the ROS from i to j is updated by the slope effect:

$$ROS_{ij} = ROS_{ij} \cdot SE_{ij}. \quad (\text{B.3})$$

Critical fireline intensity

The fireline intensity I [70] describes the rate of energy release per unit length of the fire front [kW/m], calculated by:

$$I = H \cdot W \cdot ROS \quad (\text{B.4})$$

where H represents the heat yield of the fuel [kJ/kg] and W the weight of the fuel per area [kg/m²] burned. Based on this concept, Van Wagner (1977) [352] designed a model to derive conditions for the beginning of crown fires. The model assumes that the threshold for the transition to crown fire I_{init} [kW/m] is dependent on the crown foliar moisture content FMC [%] and the crown base height (CBH) — the vertical distance between the surface and the crown.

$$I_{init} = \left(\frac{CBH (460 + 25.9FMC)}{100} \right)^{3/2} [kW/m], \quad (\text{B.5})$$

In our system, we can calculate fireline intensity I because we know the parameters H and W for each forest fuel (Canada, USA, Spain, and Chile), being able to introduce crown fire if required.

B.2 Software & Reproducibility

All code of the project can be found in the GitHub repository of this project http://www.github.com/cpaismz89/Global_Cell2Fire, as open-source software. Starting from the root directory, instances and weather scenarios can be found in the *data* folder and all code scripts are located inside the *src* folder. All C++ code for the simulator is located inside the *Cell2FireC* folder whereas the Python modules can be found inside the *DFOCell2Fire* folder.

Appendix C

Downstream Protection Value Detailed Results

C.1 Evaluation plots

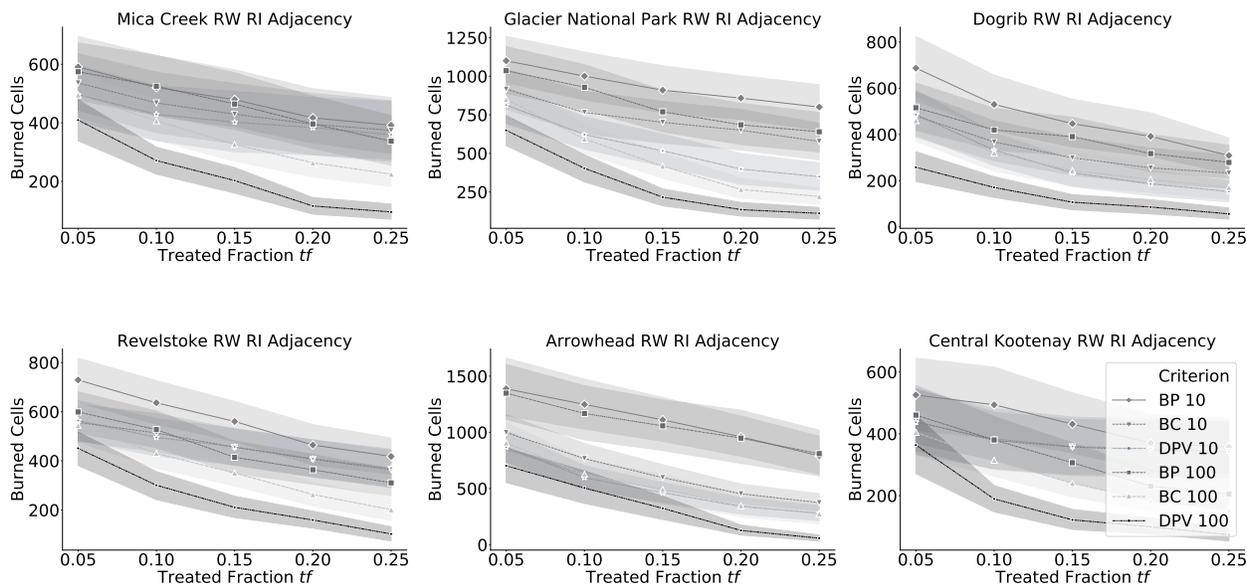


Figure C.1: Treatment plan evaluation for treatment fraction $tf \in [5\%, 25\%]$ interval. Different fraction levels are indicated in the (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications.

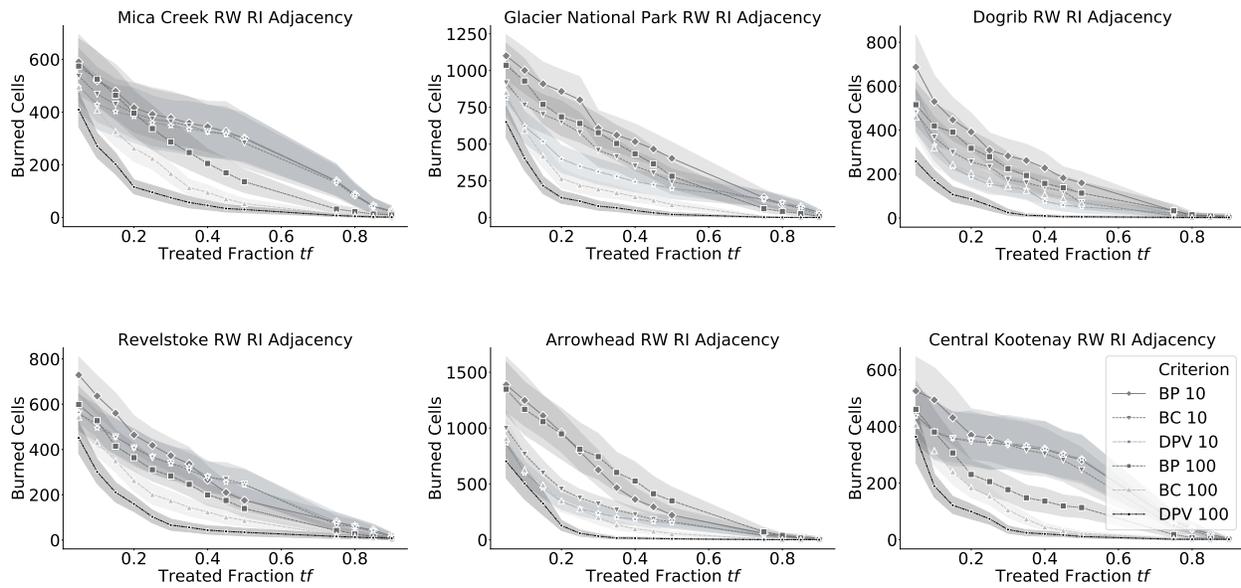


Figure C.2: Treatment plan evaluation for treatment fraction $tf \in [5\%, 100\%]$ interval. Different fraction levels are indicated in the (x-axis). We evaluate their performance using the expected area burned (y-axis). Shades along the lines associated with each plan generated from a particular fire metric trained with a certain number of replications.

C.2 Experimental Results

Instance	Trained Metric	Mean 5%-25% tf	Max	STD	$\Delta\%$ *
Sub20	BP 10	140.5	271	103.0	0
	BP 100	140.1	271	102.4	-0.3
	BC 10	104.3	271	91.6	-25.7
	BC 100	118.0	270	96.9	-16.0
	DPV 10	38.1	157	35.7	-72.9
	DPV 100	27.8	128	30.8	-80.2
Sub40	BP 10	292.0	1258	280.2	0
	BP 100	262.2	1258	256.7	-10.2
	BC 10	211.2	1208	219.2	-27.7
	BC 100	224.4	1181	227.7	-23.2
	DPV 10	196.6	1229	226.2	-32.7
	DPV 100	137.6	1211	198.2	-52.9
Sub100	BP 10	472.9	3795	560.1	0
	BP 100	384.4	2895	442.2	-18.7
	BC 10	327.0	2804	422.6	-30.9
	BC 100	279.8	2831	394.8	-40.8
	DPV 10	278.3	2417	368.6	-41.2
	DPV 100	135.7	1415	228.2	-71.3
Dogrib	BP 10	290.0	1407	327.1	0
	BP 100	263.4	1382	299.0	-9.2
	BC 10	233.4	1292	255.9	-19.5
	BC 100	125.8	806	151.9	-56.6
	DPV 10	211.8	1248	242.5	-27.0
	DPV 100	109.7	688	133.5	-62.2

Table C.1: Average number of burned cells including treatment levels $tf \in [0.05, 0.25]$ for the Alberta region instances. The proposed *DPV* model is consistently outperforming alternative metrics when trained with $|R| = 100$ replications. *Performance difference $\Delta\%$ is calculated for each instance f with respect to the worse performance — higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{f_i} - Max(Mean_f))/Max(Mean_f)$ where i is the corresponding row for the f instance.

Instance	Trained Metric	Mean 5%-25% tf	Max	STD	$\Delta\%$ *
Mica Creek	BP 10	479.8	1969	516.3	0
	BP 100	459.4	1969	500.0	-4.2
	BC 10	440.5	1953	517.0	-8.2
	BC 100	343.3	1648	330.5	-28.4
	DPV 10	414.0	1878	517.6	-13.7
	DPV 100	219.0	1595	254.4	-54.4
Central Kootenay	BP 10	435.5	2459	536.4	0
	BP 100	316.6	2047	367.8	-27.3
	BC 10	373.1	2131	490.8	-14.3
	BC 100	259.6	1637	275.7	-40.4
	DPV 10	375.0	2091	492.4	-13.9
	DPV 100	169.7	2508	280.7	-61.1
Revelstoke	BP 10	561.7	1579	444.8	0
	BP 100	443.0	1617	365.9	-21.1
	BC 10	459.3	1506	402.7	-18.2
	BC 100	357.8	1466	312.7	-36.3
	DPV 10	459.2	1487	397.2	-18.2
	DPV 100	244.9	1488	274.9	-56.4
Arrowhead	BP 10	1099.6	4666	1090.3	0
	BP 100	1066.7	4671	1191.7	-3.0
	BC 10	637.9	2686	625.9	-42.0
	BC 100	529.2	3093	691.1	-51.9
	DPV 10	510.1	2256	547.3	-53.6
	DPV 100	343.4	3054	609.0	-68.8
Glacier National Park	BP 10	934.0	3192	771.4	0
	BP 100	911.9	3107	719.2	-13.1
	BC 10	722.2	2893	691.0	-22.7
	BC 100	469.7	2424	512.0	-49.7
	DPV 10	539.9	2676	551.9	-42.2
	DPV 100	303.1	1674	412.8	-67.5

Table C.2: Average number of burned cells including treatment levels $tf \in [0.05, 0.25]$ for the British Columbia instances. The proposed *DPV* model is consistently outperforming alternative metrics when trained with $|R| = 100$ replications. *Performance difference $\Delta\%$ is calculated for each instance f with respect to the worse performance—higher average number of burned cells—obtained among all metrics. We use the formula $(Mean_{f_i} - Max(Mean_f))/Max(Mean_f)$ where i is the corresponding row for the f instance.

<i>tf</i> %	Sub20						Sub40						Sub100					
	Trained Metric	Mean	Max	STD	$\Delta\%$ *		Mean	Max	STD	$\Delta\%$ *		Mean	Max	STD	$\Delta\%$ *			
5%	BP 10	184.2	271	100.7	0		399.2	1258	308.2	0		687.0	3795	704.9	0.0			
	BP 100	181.6	271	102.1	-1.4		369.4	1258	297.1	-7.5		515.7	2895	506.4	-25.0			
	BC 10	166.8	271	103.4	-9.5		330.7	1208	274.1	-17.5		479.9	2804	480.2	-30.2			
	BC 100	169.0	270	101.9	-8.2		329.4	1181	277.9	-17.5		460.0	2831	498.4	-33.0			
	DPV 10	58.8	157	44.4	-68.1		337.5	1229	292.3	-15.4		486.0	2417	499.3	-29.3			
	DPV 100	47.3	128	39.3	-74.3		274.8	1211	275.6	-31.2		258.0	1415	334.3	-62.4			
10%	BP 10	162.6	256	102.6	0		309.5	1115	279.8	0		529.9	3089	602.6	0			
	BP 100	161.9	256	100.9	-0.5		265.7	1068	256.3	-14.1		419.4	2450	459.4	-20.9			
	BC 10	120.7	256	92.9	-25.77		256.3	1000	221.2	-17.2		367.8	2464	438.3	-30.6			
	BC 100	141.9	256	92.9	-25.8		272.8	1077	246.6	-11.9		319.7	2289	415.7	-39.7			
	DPV 10	44.8	128	37.1	-72.5		249.3	1104	250.8	-19.4		332.1	1567	371.5	-37.3			
	DPV 100	33.1	105	31.6	-79.6		186.8	856	228.4	-39.6		170.8	1144	233.2	-67.8			
15%	BP 10	137.4	240	101.4	-0.6		287.4	1073	274.6	0		446.7	2237	517.5	0			
	BP 100	138.1	240	99.7	0		247.0	1015	244.6	-14.1		390.7	2319	433.3	-12.5			
	BC 10	101.6	239	86.5	-26.4		191.5	837	195.6	-33.4		298.9	2258	407.7	-33.1			
	BC 100	120.9	240	93.5	-12.5		218.7	898	213.0	-23.9		243.5	2120	359.8	-45.5			
	DPV 10	36.6	111	31.2	-73.5		174.2	847	182.4	-39.4		234.0	1377	303.4	-47.6			
	DPV 100	24.6	81	26.3	-82.2		96.7	502	127.3	-66.4		106.9	774	173.8	-76.1			
20%	BP 10	116.5	225	97.8	-0.4		249.4	1021	262.8	0		391.8	2028	464.5	0			
	BP 100	117.0	225	97.8	0		222.2	919	229.3	-10.9		317.2	2204	399.6	-19.0			
	BC 10	81.5	223	75.0	-30.3		151.2	789	170.9	-39.4		255.4	2077	378.1	-34.8			
	BC 100	88.6	221	83.7	-24.3		162.7	778	170.8	-34.8		205.2	1958	316.4	-47.6			
	DPV 10	28.6	81	26.5	-75.5		121.6	767	146.5	-51.2		186.2	1208	276.1	-52.5			
	DPV 100	18.7	77	22.3	-84.1		71.5	471	104.9	-71.4		86.2	762	159.5	-78.0			
25%	BP 10	101.8	210	91.8	0		214.6	921	241.2	0		309.0	1916	392.2	0			
	BP 100	101.8	210	92.0	0		206.9	883	222.0	-3.6		279.1	2111	369.0	-9.7			
	BC 10	51.1	129	47.4	-49.8		126.6	742	151.6	-41.0		233.0	1899	357.5	-24.6			
	BC 100	69.4	204	71.6	-31.8		138.4	718	154.6	-35.5		170.4	1807	285.6	-44.8			
	DPV 10	21.8	77	23.8	-78.6		100.4	715	128.7	-53.2		153.6	973	237.5	-50.4			
	DPV 100	15.1	69	19.7	-85.2		58.4	447	94.1	-72.8		56.6	525	121.4	-81.7			

Table C.3: Learning capacity experiment results for Sub20, Sub40, and Sub100 instances. *Performance difference $\Delta\%$ is calculated for each treatment level *tf* with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{tf_i} - Max(Mean_{tf})) / Max(Mean_{tf})$ where *i* is the corresponding row for the *tf* level.

tf %	Trained Metric			Dogrib				Mica Creek				Central Kootenay				
	Mean	Max	STD	Mean	Max	STD	$\Delta\%$ *	Mean	Max	STD	$\Delta\%$ *	Mean	Max	STD	$\Delta\%$ *	
5%	BP 10	358.5	1407	370.4	0	590.6	1969	519.0	0	525.5	2459	605.5	0			
	BP 100	350.0	1382	359.9	-2.4	574.4	1969	520.3	-2.7	460.4	2047	503.8	-12.4			
	BC 10	325.4	1292	316.1	-9.2	536.2	1953	518.9	-9.2	439.5	2131	518.8	-16.4			
	BC 100	228.8	806	217.3	-36.2	495.1	1648	406.8	-16.2	404.9	1637	390.9	-23.0			
	DPV 10	311.4	1248	319.8	-13.2	495.8	1878	495.2	-16.0	437.4	2091	516.8	-16.8			
	DPV 100	177.0	688	166.8	-50.6	409.2	1595	340.5	-30.7	363.6	2508	494.9	-30.8			
10%	BP 10	321.3	1297	338.5	0	510.5	1956	535.4	-0.9	493.8	2352	577.9	0			
	BP 100	291.5	1224	318.4	-9.3	525.1	1969	538.6	0	380.2	1516	384.0	-23.0			
	BC 10	271.9	1095	270.0	-15.4	467.9	1884	513.9	-10.9	378.7	1961	503.3	-23.3			
	BC 100	141.4	488	145.5	-56.0	406.6	1517	362.0	-22.6	314.7	1207	280.8	-36.3			
	DPV 10	257.7	840	257.9	-19.8	429.1	1866	518.2	-18.3	383.3	1976	506.3	-22.4			
	DPV 100	124.0	497	131.9	-61.4	271.1	973	242.2	-48.4	189.5	1059	208.6	-61.6			
15%	BP 10	282.2	1189	318.1	0	479.1	1913	524.8	0	431.3	2125	522.3	0			
	BP 100	258.7	1121	290.1	-8.3	464.1	1883	515.3	-3.1	306.5	1291	312.3	-28.9			
	BC 10	227.7	882	238.4	-19.3	429.5	1853	509.4	-10.4	357.6	1810	481.2	-17.1			
	BC 100	104.3	418	119.7	-63.0	326.4	1309	294.9	-31.9	240.1	997	228.4	-44.3			
	DPV 10	199.0	730	217.2	-29.5	402.1	1856	524.3	-16.1	359.2	1821	484.6	-16.7			
	DPV 100	108.2	484	119.7	-63.0	202.8	913	213.6	-57.7	121.9	653	156.5	-71.8			
20%	BP 10	257.4	1081	309.1	0	417.0	1774	489.8	0	369.9	1805	476.8	0			
	BP 100	226.5	1000	261.5	-12.0	396.2	1705	468.4	-5.0	230.6	1025	266.2	-37.7			
	BC 10	185.4	771	211.6	-28.0	395.1	1837	517.1	-5.3	346.9	1765	475.8	-6.2			
	BC 100	90.7	346	105.4	-64.8	263.3	1155	256.9	-36.9	184.1	587	182.6	-50.2			
	DPV 10	162.4	652	188.3	-36.9	383.6	1840	526.0	-8.0	352.3	1789	480.2	-4.73			
	DPV 100	77.2	407	107.0	-70.0	116.1	555	148.0	-72.2	99.5	506	134.4	-73.1			
25%	BP 10	230.5	955	283.6	0	391.7	1774	496.5	0	357.1	1782	476.8	0			
	BP 100	190.7	898	228.5	-17.3	337.3	1545	421.4	-13.9	205.4	1025	258.8	-42.5			
	BC 10	156.5	648	193.2	-32.1	374.0	1817	519.8	-4.5	342.9	1757	476.8	-3.99			
	BC 100	64.0	249	76.6	-72.3	224.9	1002	229.9	-42.6	154.3	529	155.7	-56.8			
	DPV 10	128.4	502	152.0	-44.3	359.3	1812	523.3	-8.3	342.7	1757	476.9	-4.1			
	DPV 100	62.3	360	100.2	-73.0	95.6	550	136.0	-75.6	73.8	384	105.5	-79.3			

Table C.4: Learning capacity experiment results for Dogrib, Mica Creek, and Central Kootenay. *Performance difference $\Delta\%$ is calculated for each treatment level tf with respect to the worse performance — higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{tf_i} - Max(Mean_{tf})) / Max(Mean_{tf})$ where i is the corresponding row for the tf level.

tf %	Revelstoke						Arrowhead						Glacier National Park					
	Trained Metric		Mean	Max	STD	$\Delta\%$ *	Mean		Max	STD	$\Delta\%$ *	Mean		Max	STD	$\Delta\%$ *		
5%	BP 10	729.1	1579	466.0	0	1387.1	4666	1277.5	0	1100.3	3192	757.4	0					
	BP 100	599.0	1617	417.6	-17.9	1348.3	4671	1241.5	-2.8	1035.8	3107	777.0	-5.9					
	BC 10	556.6	1506	400.0	-23.7	997.7	2686	765.7	-28.1	916.0	2893	687.8	-16.8					
	BC 100	543.3	1466	370.1	-25.5	905.7	3093	883.6	-34.7	850.9	2424	649.0	-22.7					
	DPV 10	564.4	1487	403.0	-22.6	867.5	2256	748.1	-37.5	813.3	2676	607.7	-26.1					
	DPV 100	450.9	1488	353.8	-38.2	702.4	3054	824.5	-49.4	649.6	1674	538.0	-41.0					
10%	BP 10	636.3	1531	451.7	0	1248.5	4324	1155.2	0	1001.0	3112	772.8	0					
	BP 100	527.2	1563	383.2	-17.1	1168.0	4657	1207.3	-6.4	928.7	2881	742.9	-7.2					
	BC 10	513.3	1497	401.4	-19.3	766.4	2243	679.2	-38.6	765.4	2807	691.4	-23.5					
	BC 100	430.6	1269	309.1	-32.3	626.5	2772	709.3	-49.8	593.3	1794	519.5	-40.7					
	DPV 10	495.5	1444	402.8	-22.12	597.6	1883	509.3	-52.1	621.8	2496	552.4	-37.9					
	DPV 100	300.1	1061	280.9	-52.7	505.0	2281	728.1	-59.6	403.2	1420	441.5	-59.7					
15%	BP 10	560.5	1506	432.4	0	1111.2	3990	1048.0	0	909.8	3012	773.0	0					
	BP 100	414.1	1398	346.8	-26.0	1057.8	4604	1190.9	-4.8	769.5	2366	671.4	-15.4					
	BC 10	455.1	1483	398.2	-18.8	597.1	1877	575.3	-46.3	701.2	2752	695.7	-22.9					
	BC 100	350.8	885	274.6	-37.4	491.6	1800	453.3	-58.3	417.5	1507	419.2	-54.1					
	DPV 10	456.8	1432	392.9	-18.5	463.8	1800	453.3	-58.3	516.8	2280	528.7	-43.2					
	DPV 100	210.8	793	216.3	-62.4	324.2	1930	548.3	-70.8	216.2	1090	285.0	-76.2					
20%	BP 10	464.4	1481	412.8	0	960.1	3506	935.8	0	858.3	3012	772.1	0					
	BP 100	363.8	1232	316.4	-21.7	948.6	4540	1169.7	-1.2	685.1	2366	666.3	-20.2					
	BC 10	406.7	1437	396.2	-12.4	454.0	1776	430.0	-52.7	650.5	2541	682.3	-24.2					
	BC 100	263.1	730	245.5	-43.3	347.3	2123	516.9	-63.8	266.3	1242	322.5	-69.0					
	DPV 10	409.0	1407	382.7	-11.9	340.7	1779	378.6	-64.5	398.3	2112	486.0	-53.6					
	DPV 100	159.4	642	183.1	-65.7	127.2	805	231.0	-86.8	135.4	1053	221.9	-84.2					
25%	BP 10	418.2	1446	392.8	0	791.1	3080	911.0	-2.4	800.5	2992	760.3	0					
	BP 100	310.6	979	277.3	-25.7	810.7	4331	1096.1	0	640.1	2359	666.7	-20.0					
	BC 10	364.3	1402	395.3	-12.8	374.4	1756	396.3	-53.8	577.9	2564	663.0	-27.8					
	BC 100	201.3	698	146.6	-75.5	274.8	1937	427.9	-66.1	220.8	1150	284.8	-72.4					
	DPV 10	370.5	1366	383.2	-11.4	281.1	1737	348.0	-65.3	349.1	1919	453.3	-56.4					
	DPV 100	102.6	630	146.6	-75.5	58.4	632	139.1	-92.8	111.3	917	198.5	-86.1					

Table C.5: Learning capacity experiment results for Revelstoke, Arrowhead, and Glacier National Park. *Performance difference $\Delta\%$ is calculated for each treatment level tf with respect to the worse performance —higher average number of burned cells— obtained among all metrics. We use the formula $(Mean_{t_f,i} - Max(Mean_{t_f})) / Max(Mean_{t_f})$ where i is the corresponding row for the tf level.

Appendix D

Detailed Results of Extended Framework

D.1 Experimental results

Table D.1: Average discounted utility results for Napa instance evaluated from 100 simulations. Results for all λ combinations between the DPV heatmap and NV layers are presented by dominating utility category (column 1) and treatment fraction level (column 2).

Utility	tf	$\Delta U(0)$	$\Delta U(0.25)$	$\Delta U(0.5)$	$\Delta U(0.75)$	$\Delta U(1)$
Access	0.05	-630.36 ± 697.47	-719.95 ± 733.02	-212.01 ± 343.35	-2.29 ± 217.35	168.94 ± 123.13
	0.10	-204.23 ± 680.14	-222.66 ± 685.27	275.49 ± 260.19	444.88 ± 165.37	567.69 ± 86.51
	0.15	76.93 ± 725.53	482.08 ± 431.61	731.61 ± 220.05	800.39 ± 151.74	903.39 ± 65.42
	0.20	747.88 ± 539.28	939.52 ± 371.53	1143.22 ± 213.7	1134.43 ± 147.12	1216.88 ± 57.02
	0.25	1218.59 ± 453.58	1373.17 ± 344.61	1533.92 ± 213.11	1552.99 ± 132.53	1508.5 ± 53.91
	0.5	3298.82 ± 198.32	3324.25 ± 169.89	3423.14 ± 99.82	3352.41 ± 42.74	3097.83 ± 13.78
Carbon	0.05	-98.88 ± 264.17	-236.28 ± 369.97	-388.85 ± 298.23	-217.91 ± 196.83	-119.62 ± 146.07
	0.10	126.69 ± 258.7	15.7 ± 339.87	-155.04 ± 217.61	-8.59 ± 137.88	60.61 ± 103.34
	0.15	330.78 ± 231.39	345.71 ± 214.09	172.81 ± 178.55	154.35 ± 118.93	209.04 ± 75.63
	0.20	496.06 ± 199.97	520.36 ± 183.79	438.15 ± 161.36	404.96 ± 109.03	339.83 ± 59.88
	0.25	759.32 ± 167.12	822.89 ± 151.53	715.74 ± 139.45	608.36 ± 79.46	494.92 ± 46.97
	0.5	1661.86 ± 113.59	1623.56 ± 116.32	1559.84 ± 132.09	1533.92 ± 78.67	1170.72 ± 25.24
Equal	0.05	-393.71 ± 441.85	-471.49 ± 519.2	-346.28 ± 325.8	-139.49 ± 204.77	-1.67 ± 132.82
	0.10	-49.04 ± 430.45	-100.32 ± 480.92	19.05 ± 243.76	154.46 ± 151.34	262.15 ± 94.0
	0.15	272.19 ± 389.96	347.36 ± 302.84	356.9 ± 203.05	403.82 ± 135.65	491.04 ± 70.41
	0.20	563.05 ± 337.88	626.91 ± 260.67	692.6 ± 191.55	635.5 ± 129.05	694.99 ± 59.39
	0.25	827.96 ± 284.05	966.14 ± 236.93	1017.39 ± 186.06	958.71 ± 108.22	934.82 ± 52.73
	0.5	2373.07 ± 197.32	2251.69 ± 195.28	2237.53 ± 174.00	2288.82 ± 56.59	1980.16 ± 18.02
Forest	0.05	-29.83 ± 222.46	-151.18 ± 304.25	-367.75 ± 267.31	-267.42 ± 203.46	-181.82 ± 163.4
	0.10	196.82 ± 198.16	51.23 ± 275.67	-198.23 ± 206.55	-71.48 ± 134.11	-11.37 ± 95.93
	0.15	355.85 ± 184.95	231.59 ± 249.27	-38.34 ± 177.71	53.45 ± 107.85	105.97 ± 72.67
	0.20	523.46 ± 165.59	482.76 ± 179.07	242.87 ± 164.76	307.56 ± 107.06	229.7 ± 50.59
	0.25	756.43 ± 156.15	751.86 ± 162.06	573.52 ± 156.47	493.94 ± 78.28	350.95 ± 43.45
	0.5	1437.75 ± 119.00	1485.83 ± 83.09	1366.29 ± 102.55	1244.37 ± 69.88	957.45 ± 25.99
Population	0.05	-54.52 ± 145.67	-122.11 ± 201.67	-215.18 ± 282.34	-202.21 ± 166.13	-154.92 ± 139.03
	0.10	47.22 ± 142.73	10.4 ± 147.42	-7.44 ± 149.98	-82.36 ± 139.5	-40.6 ± 110.42
	0.15	126.16 ± 136.79	129.71 ± 112.21	99.94 ± 121.06	7.84 ± 116.75	48.0 ± 84.91
	0.20	213.52 ± 121.23	213.94 ± 107.48	172.36 ± 125.59	150.56 ± 103.08	185.98 ± 70.19
	0.25	288.26 ± 105.29	285.93 ± 118.84	241.63 ± 125.67	250.5 ± 99.98	273.0 ± 52.72
	0.5	681.34 ± 103.15	645.12 ± 127.37	856.13 ± 30.42	817.33 ± 26.46	748.70 ± 17.05

Table D.2: Average discounted utility results for the Paradise instance evaluated from 100 simulations. Results for all λ value combinations between the DPV heatmap and the NV layer are presented by dominating utility category (column 1) and treatment fraction level (column 2).

Utility	tf	$\Delta U(0)$	$\Delta U(0.25)$	$\Delta U(0.5)$	$\Delta U(0.75)$	$\Delta U(1)$
Access	0.05	178.43 \pm 318.39	219.03 \pm 277.09	271.39 \pm 171.98	314.29 \pm 106.98	351.16 \pm 69.99
	0.10	709.31 \pm 257.46	713.05 \pm 247.5	772.21 \pm 135.73	763.92 \pm 94.98	770.0 \pm 65.49
	0.15	1170.32 \pm 257.5	1182.01 \pm 235.67	1236.91 \pm 133.02	1201.7 \pm 94.7	1188.84 \pm 66.45
	0.20	1628.9 \pm 254.39	1647.97 \pm 215.35	1666.09 \pm 131.14	1639.13 \pm 95.41	1585.71 \pm 72.64
	0.25	2074.37 \pm 243.06	2096.7 \pm 200.28	2079.68 \pm 128.94	2046.8 \pm 89.94	1981.67 \pm 74.27
	0.5	4107.38 \pm 142.12	4121.74 \pm 152.96	4139.48 \pm 51.61	3974.30 \pm 15.38	3440.25 \pm 0.93
Carbon	0.05	170.0 \pm 167.37	119.33 \pm 212.13	112.99 \pm 202.85	79.86 \pm 169.74	92.99 \pm 144.35
	0.10	501.83 \pm 154.96	395.71 \pm 203.75	406.94 \pm 205.52	377.67 \pm 163.68	326.32 \pm 146.97
	0.15	880.92 \pm 119.85	829.95 \pm 143.02	749.78 \pm 139.05	621.85 \pm 66.66	594.98 \pm 62.37
	0.20	1136.25 \pm 115.12	1111.82 \pm 140.39	1000.25 \pm 82.06	851.37 \pm 35.68	770.77 \pm 30.33
	0.25	1366.23 \pm 118.15	1322.46 \pm 141.34	1267.54 \pm 79.92	1079.41 \pm 23.32	930.57 \pm 13.21
	0.5	2325.86 \pm 75.56	2263.15 \pm 73.76	2242.70 \pm 17.81	2128.67 \pm 8.69	1917.48 \pm 1.24
Equal	0.05	146.99 \pm 253.0	96.19 \pm 298.17	126.65 \pm 192.1	139.03 \pm 133.54	163.79 \pm 96.69
	0.10	517.53 \pm 254.44	463.82 \pm 299.69	454.8 \pm 188.2	430.92 \pm 131.2	402.6 \pm 97.28
	0.15	869.89 \pm 250.2	814.78 \pm 293.29	846.11 \pm 114.27	716.38 \pm 133.36	658.84 \pm 98.7
	0.20	1338.95 \pm 210.91	1296.81 \pm 237.16	1148.0 \pm 85.11	1026.82 \pm 31.18	906.99 \pm 103.73
	0.25	1740.49 \pm 207.27	1744.59 \pm 192.15	1516.46 \pm 71.49	1302.24 \pm 21.32	1154.24 \pm 107.06
	0.5	3295.83 \pm 103.08	3286.24 \pm 104.47	3190.06 \pm 45.89	2953.84 \pm 13.94	2390.28 \pm 1.29
Forest	0.05	216.81 \pm 202.38	249.91 \pm 204.67	213.87 \pm 184.86	198.49 \pm 157.21	190.93 \pm 134.61
	0.10	683.54 \pm 154.58	694.02 \pm 161.88	624.48 \pm 165.31	545.56 \pm 150.06	479.31 \pm 137.23
	0.15	1024.93 \pm 150.04	1017.27 \pm 141.46	990.61 \pm 110.4	848.03 \pm 62.04	778.14 \pm 56.25
	0.20	1298.57 \pm 148.71	1296.31 \pm 142.27	1250.72 \pm 69.61	1115.27 \pm 37.33	991.0 \pm 28.6
	0.25	1528.34 \pm 140.52	1562.9 \pm 106.9	1500.39 \pm 60.08	1395.67 \pm 31.94	1192.01 \pm 12.08
	0.5	2598.65 \pm 64.54	2566.28 \pm 63.29	2551.14 \pm 17.53	2452.89 \pm 8.03	2296.49 \pm 0.90
Population	0.05	324.29 \pm 70.44	268.57 \pm 125.4	-26.16 \pm 106.26	-12.03 \pm 91.99	-0.25 \pm 86.53
	0.10	695.32 \pm 71.5	630.53 \pm 129.06	354.47 \pm 104.33	91.32 \pm 87.87	93.99 \pm 82.53
	0.15	955.1 \pm 73.52	889.47 \pm 127.92	707.54 \pm 105.29	183.65 \pm 88.69	173.49 \pm 85.17
	0.20	1164.93 \pm 72.59	1101.32 \pm 123.66	928.23 \pm 102.43	418.05 \pm 93.16	241.01 \pm 90.85
	0.25	1333.8 \pm 69.21	1279.45 \pm 69.42	1106.89 \pm 103.3	811.5 \pm 93.61	319.53 \pm 93.13
	0.5	2007.73 \pm 35.38	1958.83 \pm 62.23	1852.77 \pm 114.98	1741.65 \pm 78.41	708.24 \pm 3.74

Table D.3: Average discounted utility results for the Getty center instance evaluated from 100 simulations. Results for all λ value combinations between the DPV heatmap and the NV layer are presented by dominating utility category (column 1) and treatment fraction level (column 2).

Utility	tf	$\Delta U(0)$	$\Delta U(0.25)$	$\Delta U(0.5)$	$\Delta U(0.75)$	$\Delta U(1)$
Access	0.05	-472.96 ± 453.7	-392.1 ± 417.92	-262.43 ± 303.32	-139.59 ± 230.58	-80.92 ± 201.81
	0.10	-353.71 ± 443.07	-227.01 ± 375.65	-72.71 ± 266.06	4.14 ± 189.54	69.61 ± 150.55
	0.15	-205.64 ± 434.12	-75.1 ± 364.49	82.81 ± 213.83	160.22 ± 136.52	211.05 ± 103.45
	0.20	-64.41 ± 431.02	82.69 ± 350.25	225.15 ± 184.66	284.95 ± 120.86	330.84 ± 74.25
	0.25	125.28 ± 407.03	211.91 ± 338.75	355.09 ± 159.94	406.88 ± 97.7	445.96 ± 51.63
	0.5	905.18 ± 291.65	949.64 ± 256.68	971.89 ± 138.03	967.81 ± 70.83	988.54 ± 24.95
Carbon	0.05	-205.02 ± 215.85	-252.35 ± 283.41	-250.66 ± 255.97	-177.37 ± 220.13	-151.03 ± 213.97
	0.10	-95.77 ± 205.28	-131.89 ± 259.1	-85.9 ± 175.91	-83.15 ± 171.28	-47.19 ± 142.25
	0.15	10.27 ± 191.77	-48.71 ± 245.27	7.7 ± 146.14	33.27 ± 111.25	50.12 ± 103.33
	0.20	106.72 ± 176.78	51.74 ± 222.78	112.5 ± 140.3	104.59 ± 82.4	114.97 ± 82.65
	0.25	183.74 ± 174.64	121.02 ± 217.3	187.24 ± 118.38	170.12 ± 69.47	178.04 ± 51.78
	0.5	578.29 ± 149.24	569.52 ± 136.31	514.78 ± 94.54	492.50 ± 48.50	467.99 ± 28.17
Equal	0.05	-291.81 ± 270.13	-340.11 ± 348.64	-267.66 ± 274.53	-156.46 ± 214.39	-121.35 ± 200.53
	0.10	-182.42 ± 264.34	-238.24 ± 329.64	-127.19 ± 229.6	-52.0 ± 168.74	-18.22 ± 153.08
	0.15	-77.81 ± 262.55	-137.06 ± 321.14	10.99 ± 165.64	64.96 ± 117.05	82.66 ± 107.17
	0.20	11.86 ± 255.07	-37.58 ± 311.16	96.22 ± 150.43	146.98 ± 100.09	170.95 ± 82.78
	0.25	96.76 ± 246.77	31.2 ± 307.0	203.37 ± 136.78	225.95 ± 82.37	244.5 ± 52.99
	0.5	636.70 ± 186.57	633.69 ± 165.57	589.46 ± 112.89	600.49 ± 54.09	599.59 ± 26.80
Forest	0.05	-154.77 ± 168.04	-216.76 ± 232.17	-285.57 ± 288.36	-226.0 ± 244.0	-164.97 ± 214.63
	0.10	-56.82 ± 156.05	-114.55 ± 208.5	-133.43 ± 205.9	-115.09 ± 178.6	-72.47 ± 144.47
	0.15	19.91 ± 144.58	-26.92 ± 189.24	-45.63 ± 158.45	-11.54 ± 119.47	15.09 ± 103.73
	0.20	96.95 ± 133.05	36.43 ± 176.55	38.2 ± 128.73	48.1 ± 96.83	67.37 ± 77.59
	0.25	147.82 ± 131.6	110.52 ± 165.47	115.68 ± 120.95	108.12 ± 77.0	125.16 ± 53.59
	0.5	474.79 ± 99.54	474.01 ± 102.52	390.05 ± 94.46	369.58 ± 54.05	359.10 ± 28.73
Population	0.05	-70.86 ± 65.58	-112.05 ± 109.7	-193.76 ± 204.4	-173.36 ± 196.66	-170.81 ± 198.96
	0.10	-43.46 ± 64.26	-70.31 ± 89.42	-136.47 ± 160.56	-115.5 ± 150.06	-112.4 ± 151.19
	0.15	-16.44 ± 64.65	-27.52 ± 66.72	-75.13 ± 113.62	-59.41 ± 105.36	-56.56 ± 105.75
	0.20	7.38 ± 61.85	-4.0 ± 59.09	-38.04 ± 93.7	-21.94 ± 85.87	-19.29 ± 84.26
	0.25	35.15 ± 59.66	22.34 ± 55.45	-2.93 ± 74.14	14.77 ± 55.57	18.24 ± 53.49
	0.5	158.55 ± 43.87	123.64 ± 43.91	106.78 ± 48.59	119.83 ± 32.95	123.32 ± 28.41

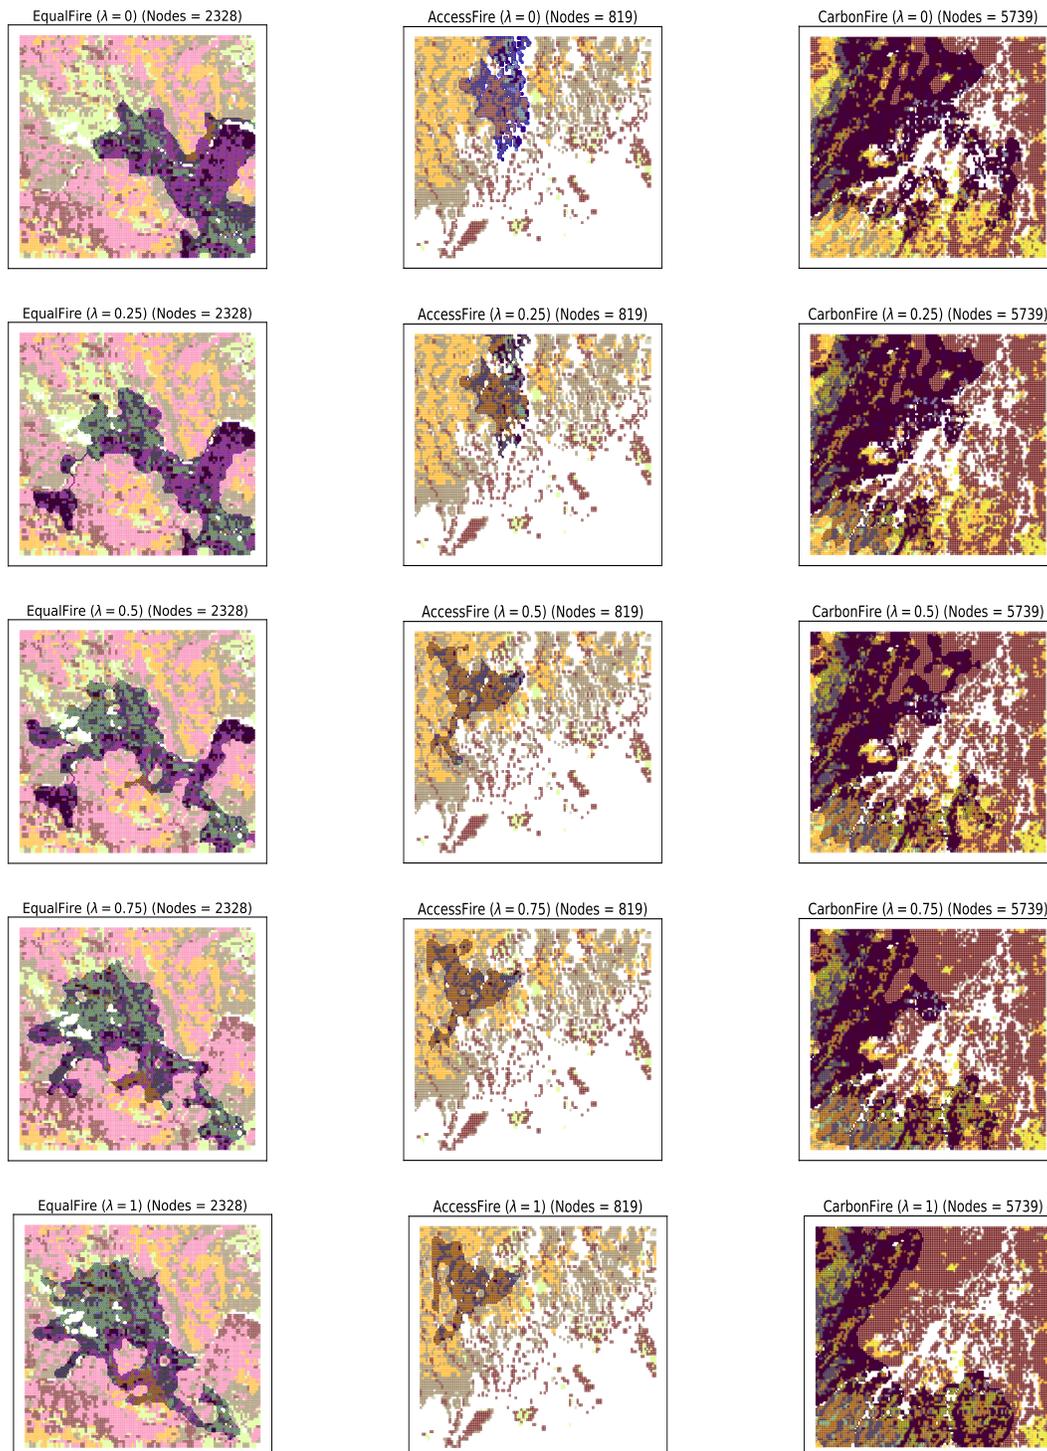


Figure D.1: Resource allocation sample plans for Napa valley, Getty center, and Paradise (columns) instances for different λ weights (rows) to account for the expected losses due to future wildfires, at a specific treatment fraction tf . Significant variations in the optimal plans can be observed as the λ values are modified to include future wildfire risk into the objective function. Original land cover colors have been modified for better contrast and non-flammable nodes have been removed (white space).

D.2 Availability of data and material

All data products and sub-products have been recorded and organized into a series of GeoTif files, ready to share with other researchers and the community in an open-source repository (http://www.github.com/cpaismz89/DPV_Utility_Extension) and Google Earth Engine assets. The code to reproduce every step of our research is available as a series of Python scripts and Jupyter notebooks for visualization convenience.

Appendix E

Delineating Fire Regimes with AI

E.1 Framework and clustering results

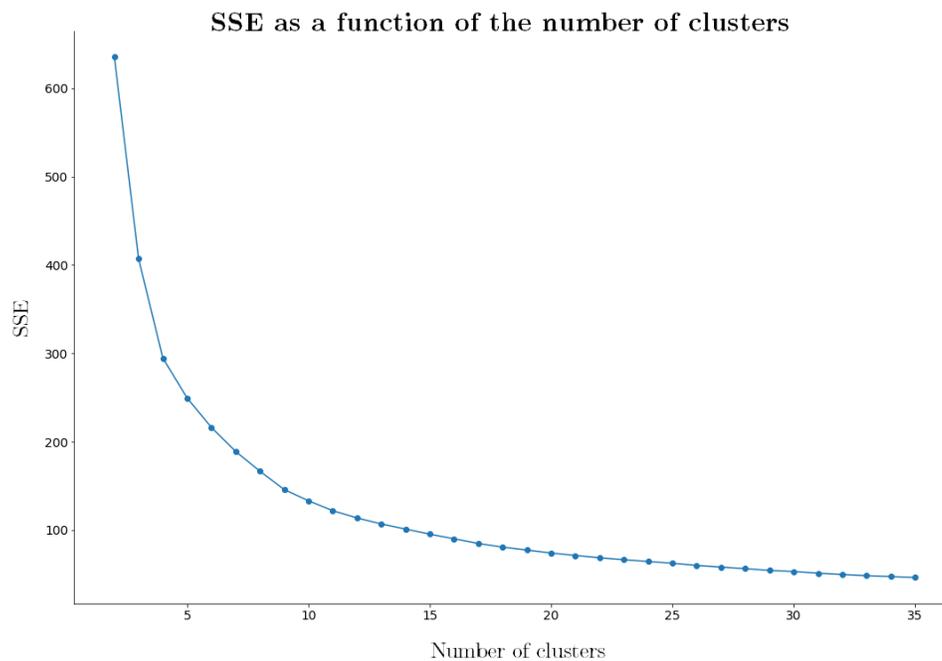


Figure E.1: Cluster convergence. Sum of squared distances of samples with respect to the nearest cluster center as a function of the number of clusters k using the K-means algorithm on the data obtained after training our self organizing map. As expected, larger values of k lead to lower SSE values, converging towards 0. We found significant variations in the slope of the function in the [15,20] interval across all tested algorithms.

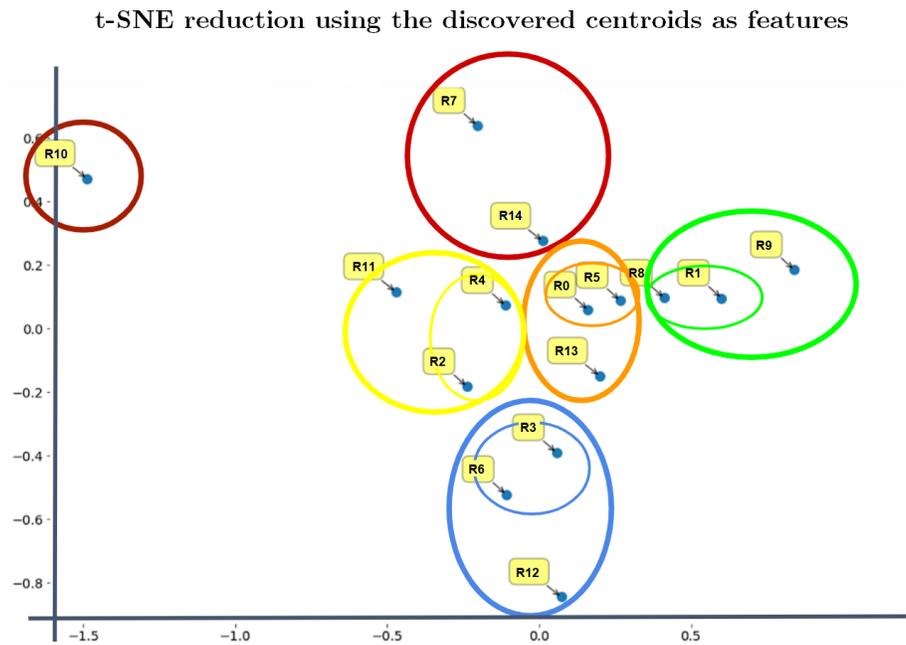


Figure E.2: Dimensionality reduction. Two-dimensional reduction using the t-SNE algorithm with the centroids of the discovered 15 regimes. From the plot, it is possible to observe the clear six macro-groups (highlighted with ovals of multiple colors) and the differences between the regimes.

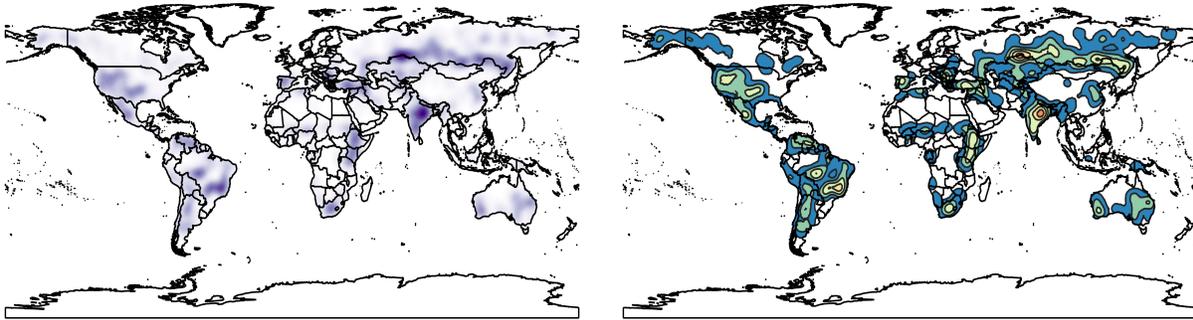


Figure E.3: Pyrome 0. R0 spatial distribution (left) and hot spots (right) representing local fire regimes.

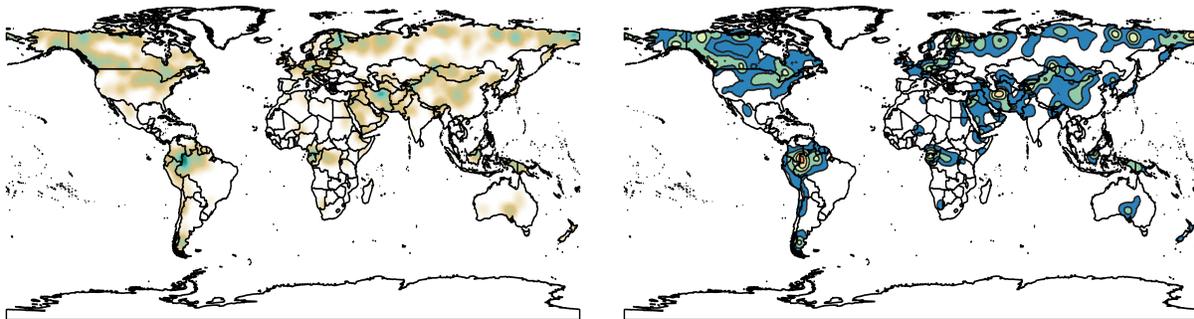


Figure E.4: Pyrome 1. R1 spatial distribution (left) and hot spots (right) representing local fire regimes.

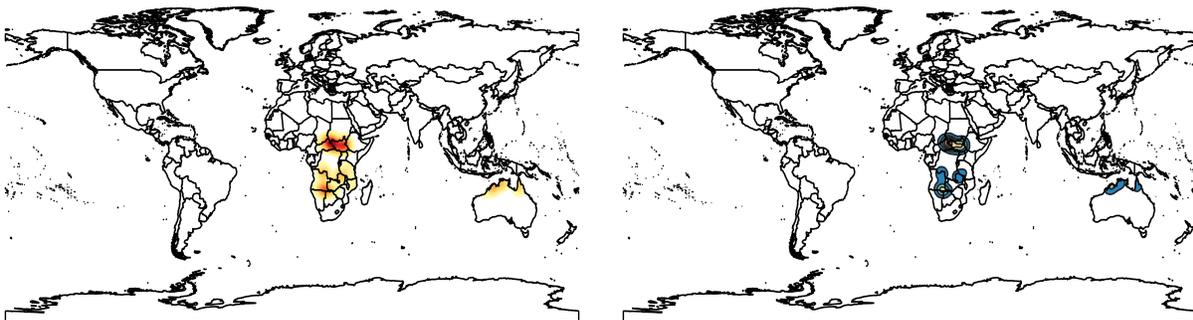


Figure E.5: Pyrome 2. R2 spatial distribution (left) and hot spots (right) representing local fire regimes.

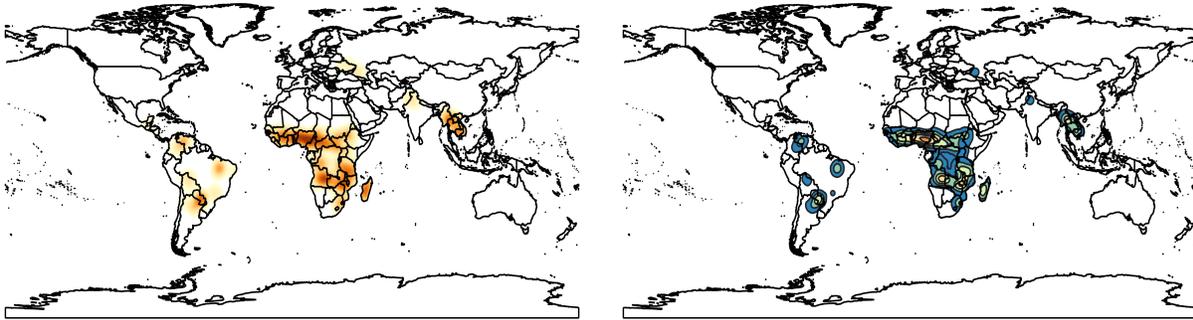


Figure E.6: Pyrome 3. R3 spatial distribution (left) and hot spots (right) representing local fire regimes.

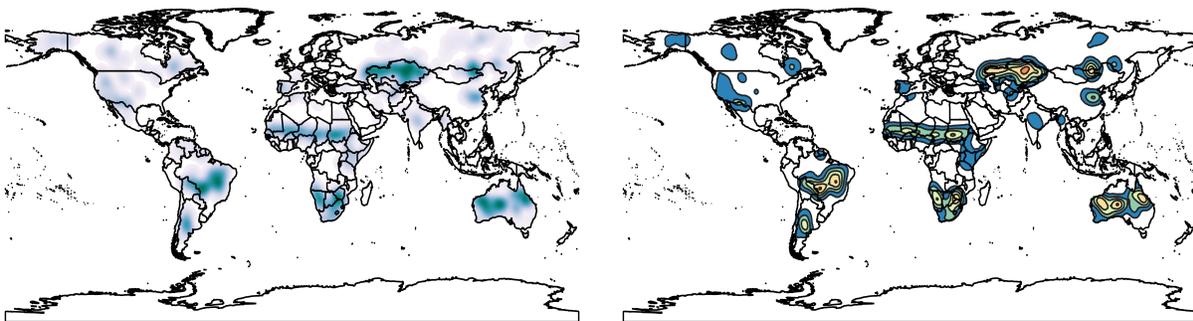


Figure E.7: Pyrome 4. R4 spatial distribution (left) and hot spots (right) representing local fire regimes.

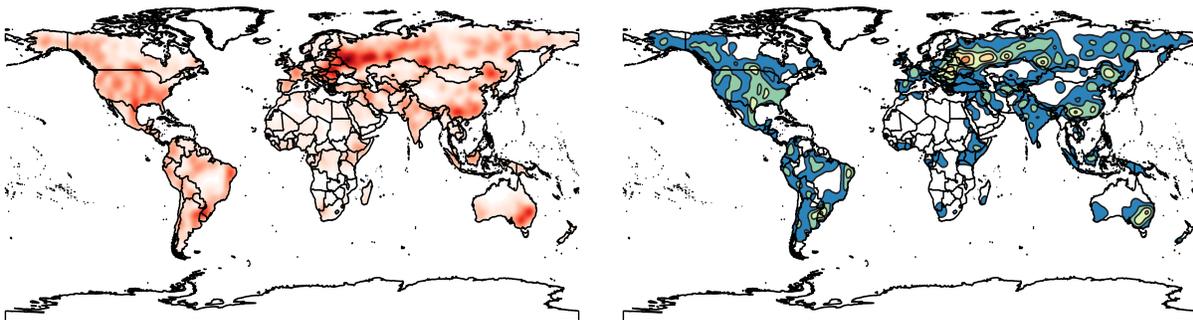


Figure E.8: Pyrome 5. R5 spatial distribution (left) and hot spots (right) representing local fire regimes.

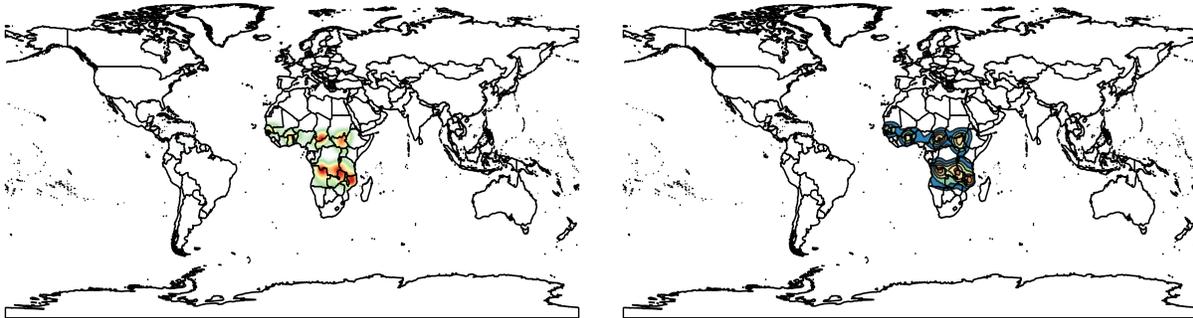


Figure E.9: Pyrome 6. R6 spatial distribution (left) and hot spots (right) representing local fire regimes.

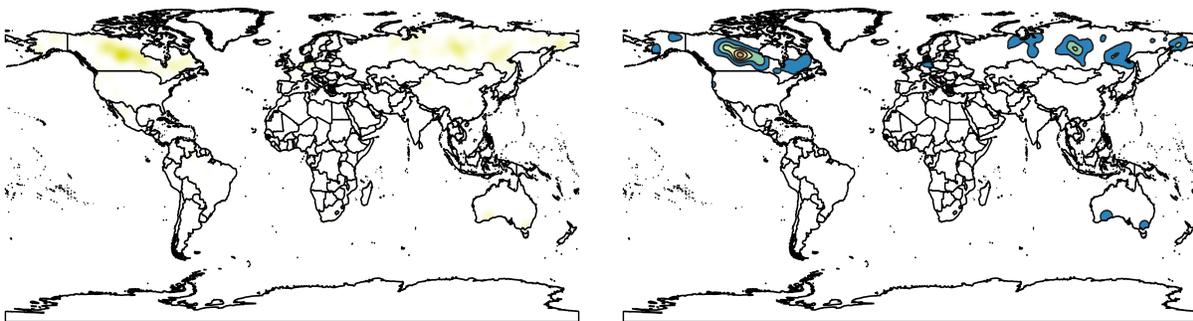


Figure E.10: Pyrome 7. R7 spatial distribution (left) and hot spots (right) representing local fire regimes.

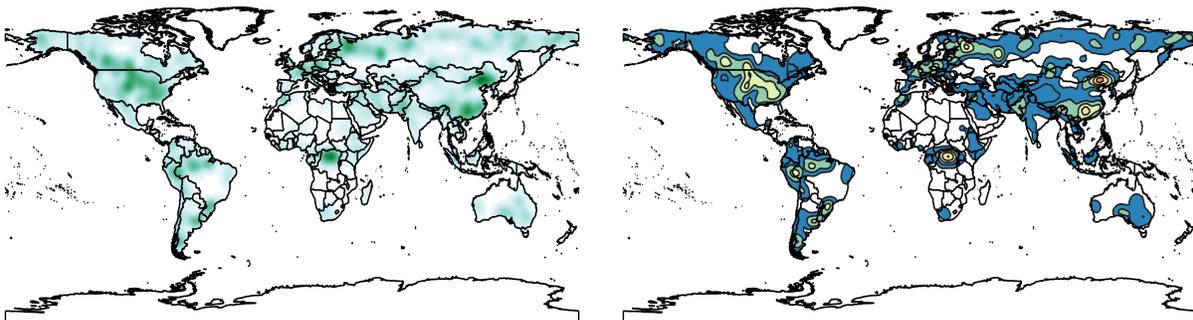


Figure E.11: Pyrome 8. R8 spatial distribution (left) and hot spots (right) representing local fire regimes.

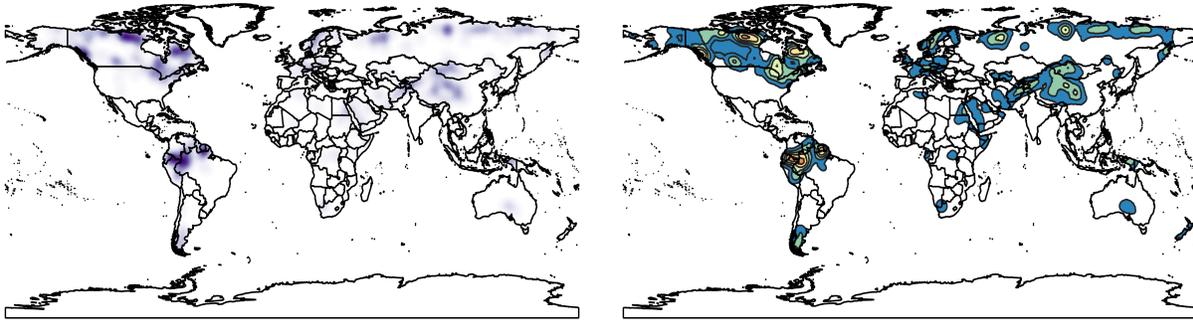


Figure E.12: Pyrome 9. R9 spatial distribution (left) and hot spots (right) representing local fire regimes.

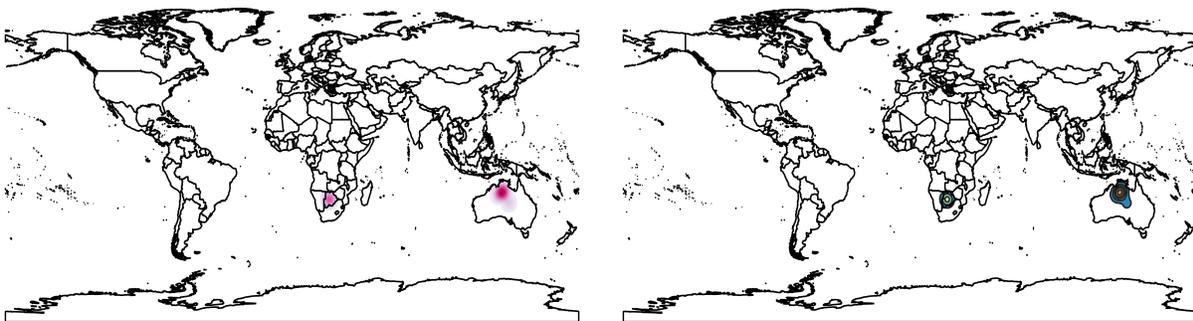


Figure E.13: Pyrome 10. R10 spatial distribution (left) and hot spots (right) representing local fire regimes.

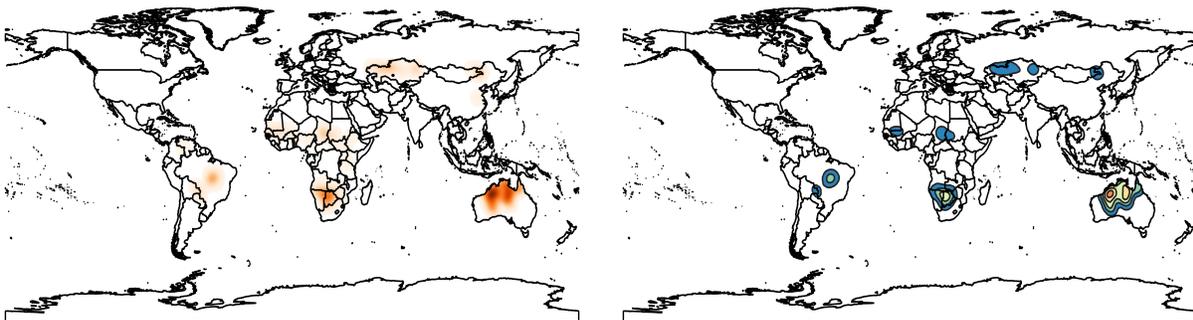


Figure E.14: Pyrome 11. R11 spatial distribution (left) and hot spots (right) representing local fire regimes.

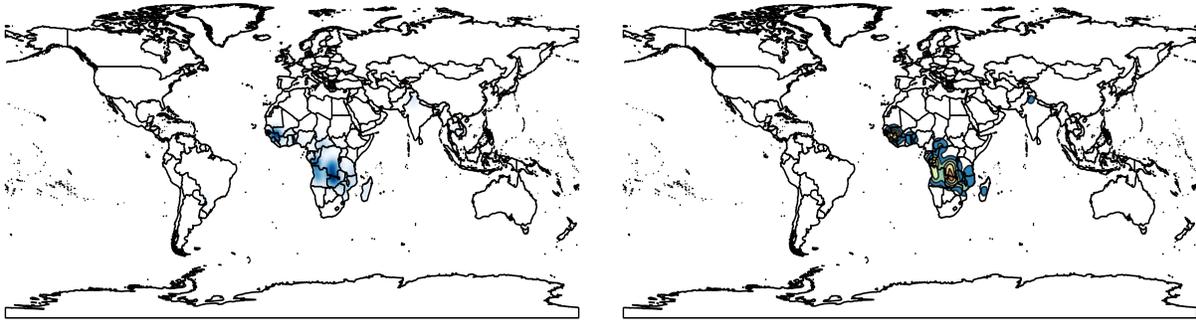


Figure E.15: Pyrome 12. R12 spatial distribution (left) and hot spots (right) representing local fire regimes.

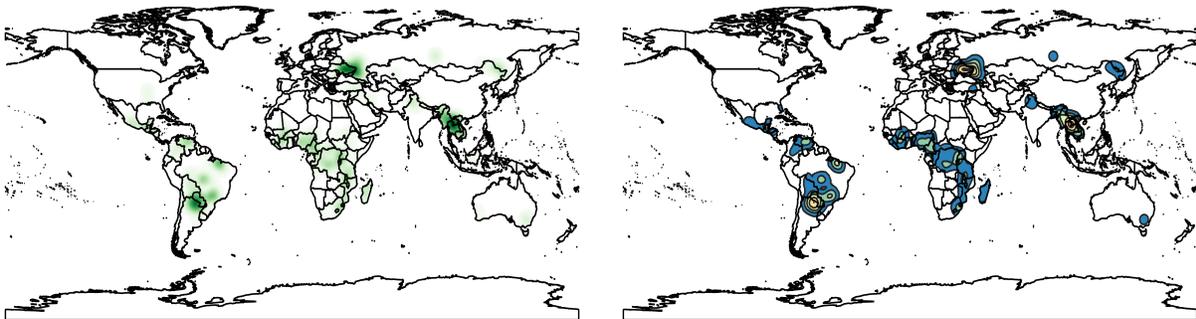


Figure E.16: Pyrome 13. R13 spatial distribution (left) and hot spots (right) representing local fire regimes.

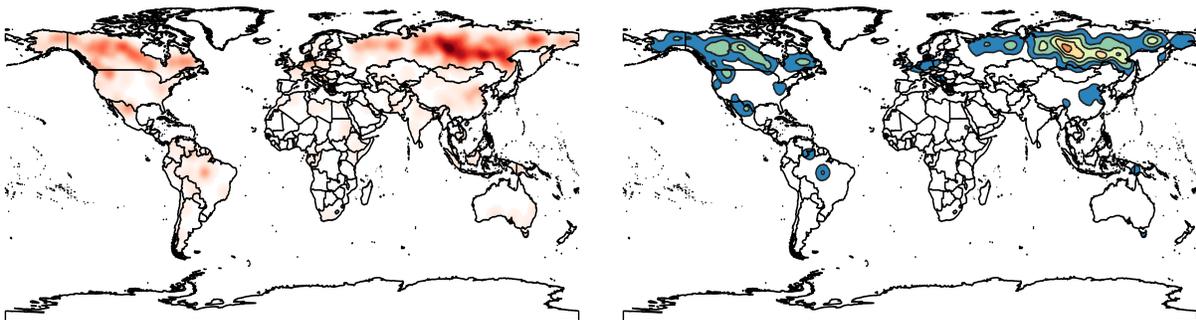


Figure E.17: Pyrome 14. R14 spatial distribution (left) and hot spots (right) representing local fire regimes.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
RO	AVG Frequency	31.89	AVG PDSI	AVG GDP [US dollars]	17,929.78 ± 18,686.63
			AVG Water deficit [mm]	AVG Population density [ppt/km ²]	GRS 51.5%
			AVG Temperature [K°]	AVG Accessability [min]	CRO 26.7%
	AVG # of Fires	50,931	AVG Max temperature [K°]		NV 10.3%
			AVG Precipitation [m]		OSL 4.7%
	AVG Size	5.22	AVG PDSI	AVG GDP [US dollars]	36,757.34 ± 11,899.11
			AVG Water deficit [mm]	AVG Population density [ppt/km ²]	OSL 13.8%
			AVG Temperature [K°]	AVG Accessability [min]	SAV 8.2%
	AVG Perimeter	9.66	AVG Max temperature [K°]		CRO 7.1%
			AVG Precipitation [m]		WDS 6.3%
			AVG PDSI	AVG GDP [US dollars]	10,103.61 ± 5149.5
	AVG Duration	4.76	AVG Water deficit [mm]	AVG Population density [ppt/km ²]	GRS 51.5%
			AVG Temperature [K°]	AVG Accessability [min]	CRO 26.7%
			AVG Max temperature [K°]		NV 10.3%
	AVG Expansion	0.77	AVG Precipitation [m]		OSL 4.7%
		AVG PDSI	AVG GDP [US dollars]	2,959.19 ± 1,179.19	
		AVG Water deficit [mm]	AVG Population density [ppt/km ²]	CRO 77.2%	
AVG Perimeter/Area	3.06	AVG Temperature [K°]	AVG Accessability [min]	GRS 6.6%	
		AVG Max temperature [K°]		WDS 4.9%	
		AVG Precipitation [m]		SAV 4.4%	
		AVG PDSI	AVG GDP [US dollars]	MFS 2.4%	
N° of cells (res 1°)	2,057	Area 2,160,000 km ²	AVG Population density [ppt/km ²]	1,285.5 ± 685.69	
		Area 4,004,000 km ²	AVG Accessability [min]	GRS 61.4%	
		Area 2,297,000 km ²		CRO 9.6%	
		Area 2,125,000 km ²	AVG GDP [US dollars]	SAV 8.8%	
		Area 2,006,000 km ²	AVG Population density [ppt/km ²]	NV 7.4%	
Total # of fires	814,896		AVG Accessability [min]	OSL 6.1%	

Table E.1: Pyromes and fire regimes details. Tables E.1-E.15 provide a comprehensive description of all pyromes and regimes. Pyromes are characterized using the inter-annual averages of fire behavior features including frequency, the number of fires, wildfire size [km^2], perimeter [km], duration [days], daily expansion [km^2/day], perimeter-per-area ratio; and the total number of cells and fires. Regimes within the 30% densest areas are characterized by their (1) spatial location; (2) average climatic conditions considering Palmer drought severity index (PDSI), water deficit [mm], temperature [K°], max temperature [K°], and total precipitation [m]; and (3) average socio-economic descriptors including the gross domestic product (GDP) [US\$], population density (population per [km^2]), accessibility (travel time [min]) to the nearest densely-populated areas with 1,500 or more inhabitants per km^2), and land-use configuration. Land use includes the following categories: closed shrublands (CSL), croplands (CRO), deciduous broadleaf forests (DBF), evergreen broadleaf palmate (EBP), evergreen needleleaf conifer (ENC), grasslands (GRS), mixed forest (MFS), non-vegetated (NV), open shrublands (OSL), permanent wetlands (PWL), savannas (SAV), water bodies (WBS), and woody savannas (WDS).

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	1.63	AVG PDSI -44.22 ± 112.27	AVG GDP [US dollars] 42,948.77 ± 9,099.68	WDS 22.2%
	AVG # of Fires	5,324.68	AVG Water deficit [mm] 123.02 ± 169.04	AVG Population density [pp]/ km ² 2.90 ± 33.48	GRS 15.6%
	AVG Size	0.79	AVG Temperature [K] 271.34 ± 11.07	AVG Accessability [min] 1,266.64 ± 1,118.64	ENC 15.1%
	AVG Perimeter	3.66	AVG Max temperature [K] 284.45 ± 10.41	AVG GDP [US dollars] 9,339.71 ± 2,859.48	OSL 11.6%
			AVG Precipitation [mm] 0.07 ± 0.02	AVG Population density [pp]/ km ² 3.31 ± 22.08	EEP 95.3%
			AVG PDSI 75.39 ± 159.71	AVG Accessability [min] 2,327.39 ± 1,471.56	SAV 2.7%
			AVG Water deficit [mm] 51.76 ± 56.7		
			AVG Temperature [K] 298.62 ± 0.61		
			AVG Max temperature [K] 306.29 ± 0.92		
			AVG Precipitation [mm] 0.23 ± 0.06		
R1	AVG Duration	2.70	AVG PDSI 82.67 ± 172.89	AVG GDP [US dollars] 41,295.28 ± 4,414.2	CHO 35.9%
	AVG Expansion	0.30	AVG Water deficit [mm] 16.08 ± 164.43	AVG Population density [pp]/ km ² 37.72 ± 147.89	MFS 16.8%
	AVG Perimeter/Area	6.59	AVG Temperature [K] 279.64 ± 10.86	AVG Accessability [min] 334.84 ± 616.89	WDS 13.9%
			AVG Max temperature [K] 283.64 ± 9.69	AVG GDP [US dollars] 37,944.20 ± 9,218.88	WBS 11.1%
			AVG Precipitation [mm] 0.08 ± 0.02		
			AVG PDSI -116.55 ± 236.67		
			AVG Water deficit [mm] 47.38 ± 116.56		
			AVG Temperature [K] 264.27 ± 15.14		
			AVG Max temperature [K] 278.33 ± 13.26		
			AVG Precipitation [mm] 0.04 ± 0.02		
			AVG PDSI 19.34 ± 177.1		
			AVG Water deficit [mm] 415.79 ± 388.96		
			AVG Temperature [K] 274.29 ± 11.54		
			AVG Max temperature [K] 286.77 ± 11.25		
			AVG Precipitation [mm] 0.05 ± 0.02		
	N° of cells (res 1°)	1,335	AVG PDSI 19.34 ± 177.1	AVG GDP [US dollars] 5,244.49 ± 4,634.23	GRS 53%
	Total # of fires	85,195	AVG Water deficit [mm] 415.79 ± 388.96	AVG Population density [pp]/ km ² 21.17 ± 97.56	NV 33.5%
			AVG Temperature [K] 274.29 ± 11.54		MFS 4.6%
			AVG Max temperature [K] 286.77 ± 11.25		CHO 3.8%
			AVG Precipitation [mm] 0.05 ± 0.02		WDS 2.5%

Table E.2: Pyrome 1. R1 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	398.99	AVG PDSI -213.83 ± 220.22	AVG GDP [US dollars] 1,928.77 ± 872.34	SAV 43.4%
	AVG # of Fires	64,393.3	AVG Water deficit [mm] 777.56 ± 603.69	AVG Population density [pp]/ km ² 12.75 ± 18.15	GRS 41.8%
	AVG Size	34.03	AVG Temperature [K] 300.57 ± 1.86	AVG Accessability [min] 500.11 ± 377.59	EBP 3.2%
	AVG Perimeter	24.22	AVG Max temperature [K] 309.19 ± 2.8		MFS 3.2%
	AVG Duration	6.24	AVG Precipitation [mm] 0.08 ± 0.07	AVG GDP [US dollars] 6,029.77 ± 2,413.07	WDS 3.1%
R2	AVG Expansion	2.28	AVG PDSI 244.03 ± 261.56	AVG Population density [pp]/ km ² 2.75 ± 5.16	GRS 74.7%
	AVG Perimeter/Area	0.94	AVG Water deficit [mm] 915.59 ± 621.05		SAV 9%
	N° of cells (res 1°)	93	AVG Temperature [K] 295.61 ± 3.04	AVG Accessability [min] 587.89 ± 273.43	OSL 6.8%
	Total # of fires	965,900	AVG Max temperature [K] 305.98 ± 2.95		CSL 4.2%
			AVG Precipitation [mm] 0.07 ± 0.09		WDS 3.1%

Table E.3: Pyrome 2. R2 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	616.62	AVG PDSI -106.97 ± 152.14 AVG Water deficit [mm] 676.12 ± 518.87 AVG Temperature [K] 300.1 ± 1.72 AVG Max temperature [K] 308.7 ± 2.48 AVG Precipitation [mm] 0.09 ± 0.07	AVG GDP [US dollars] 2,035.66 ± 1,089.43 AVG Population density [ppl/km2] 57.9 ± 153.67 AVG Accessibility [min] 256.18 ± 292.03	GRS 36.3% SAV 35% CRO 12.6% EBP 7.7% WDS 3.4%
	AVG # of Fires	398,406	Area 4,202,000 km ²		
	AVG Size	3.37	R3-a Area 4,202,000 km ²		
	AVG Perimeter	7.79	R3-b Area 3,176,000 km ²		
	AVG Duration	4.30	R3-c Area 960,000 km ²		
	AVG Expansion	0.58	R3-d Area 531,000 km ²		
	AVG Perimeter/Area	2.72	R3-e Area 463,000 km ²		
	N° of cells (res P)	333			
	Total # of fires	6,374,490			
			AVG PDSI -105.18 ± 114.28 AVG Water deficit [mm] 585.16 ± 421.72 AVG Temperature [K] 295.98 ± 2.03 AVG Max temperature [K] 304.34 ± 1.95 AVG Precipitation [mm] 0.12 ± 0.13	AVG GDP [US dollars] 7,532.68 ± 2,720.27 AVG Population density [ppl/km2] 17.3 ± 79.8 AVG Accessibility [min] 170.68 ± 129.09	SAV 63.6% WDS 12.9% GRS 11.9% CRO 10.5%
			AVG PDSI -95.69 ± 248.4 AVG Water deficit [mm] 329.22 ± 326.19 AVG Temperature [K] 298.61 ± 2.07 AVG Max temperature [K] 307.02 ± 2.09 AVG Precipitation [mm] 0.15 ± 0.12	AVG GDP [US dollars] 3,480.65 ± 3,402.52 AVG Population density [ppl/km2] 86.35 ± 125.48 AVG Accessibility [min] 180.41 ± 154.04	SAV 45.5% GRS 25.7% WDS 15.8% DRF 3.1% MFS 2.8%
			AVG PDSI -195.18 ± 114.28 AVG Water deficit [mm] 585.16 ± 421.72 AVG Temperature [K] 295.98 ± 2.03 AVG Max temperature [K] 304.34 ± 1.95 AVG Precipitation [mm] 0.12 ± 0.13	AVG GDP [US dollars] 1,420.92 ± 0.0 AVG Population density [ppl/km2] 31.96 ± 88.41 AVG Accessibility [min] 486.82 ± 265.72	GRS 76.4% WDS 8.2% SAV 6.4% EBP 5.2%

Table E.4: Pyrome 3. R3 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %		
	AVG Frequency	21.36	AVG PDSI	-469.13 ± 89.58	AVG GDP [US dollars]	29,667.85 ± 33,568.92	MFS 30.3%
	AVG # of Fires	33,982.25	AVG Water deficit [mm]	130.74 ± 170.44	AVG Population density [pop./km2]	43.7 ± 196.36	CHO 28.5%
			AVG Temperature [K°]	277.5 ± 10.69			SAV 12.5%
			AVG Max temperature [K°]	289.83 ± 11.0			WDS 9.5%
			AVG Precipitation [in]	0.06 ± 0.02			PWL 4.4%
	AVG Size	2.58	AVG PDSI	64.43 ± 144.38	AVG GDP [US dollars]	38,280.27 ± 9,084.44	CHO 33.7%
			AVG Water deficit [mm]	361.25 ± 224.51	AVG Population density [pop./km2]	37.34 ± 181.84	GRS 24.6%
			AVG Temperature [K°]	286.85 ± 8.67			WDS 20.7%
			AVG Max temperature [K°]	300.43 ± 6.74	AVG Accessibility [min]	98.03 ± 78.04	DBF 7.9%
	AVG Perimeter	6.92	AVG Precipitation [in]	0.07 ± 0.02	AVG GDP [US dollars]	4,139.49 ± 1,723.21	SAV 7.6%
			AVG PDSI	-156.3 ± 155.6	AVG Population density [pop./km2]	194.05 ± 322.13	WDS 38.4%
			AVG Water deficit [mm]	140.68 ± 96.48			SAV 31.7%
			AVG Temperature [K°]	289.24 ± 6.43	AVG Accessibility [min]	207.65 ± 248.61	GRS 5.8%
			AVG Max temperature [K°]	299.53 ± 5.29			CHO 5.4%
			AVG Precipitation [in]	0.14 ± 0.08	AVG GDP [US dollars]	41,856.29 ± 5,800	GRS 56.8%
	AVG Expansion	0.45	AVG PDSI	-71.78 ± 122.06	AVG Population density [pop./km2]	8.71 ± 58	OSL 11.8%
			AVG Water deficit [mm]	651.34 ± 558.56			WDS 10.2%
			AVG Temperature [K°]	281.98 ± 8.85	AVG Accessibility [min]	185.27 ± 97.05	ENC 9.6%
			AVG Max temperature [K°]	295.39 ± 9.38			NV 5%
			AVG Precipitation [in]	0.04 ± 0.02	AVG GDP [US dollars]	35,297.79 ± 1661.34	GRS 38.7%
	AVG Perimeter/Area	4.15	AVG PDSI	-65.74 ± 216.61	AVG Population density [pop./km2]	2.62 ± 18.13	OSL 36.3%
			AVG Water deficit [mm]	985.2 ± 595.58	AVG Accessibility [min]	323.16 ± 193.05	CHO 10.4%
			AVG Temperature [K°]	291.46 ± 5.78			SAV 7.5%
			AVG Max temperature [K°]	303.02 ± 7.02			WDS 2.5%
			AVG Precipitation [in]	0.04 ± 0.02			
	N° of cells (res 1°)	2,735	AVG PDSI	-65.74 ± 216.61	AVG GDP [US dollars]	35,297.79 ± 1661.34	GRS 38.7%
			AVG Water deficit [mm]	985.2 ± 595.58	AVG Population density [pop./km2]	2.62 ± 18.13	OSL 36.3%
			AVG Temperature [K°]	291.46 ± 5.78	AVG Accessibility [min]	323.16 ± 193.05	CHO 10.4%
			AVG Max temperature [K°]	303.02 ± 7.02			SAV 7.5%
			AVG Precipitation [in]	0.04 ± 0.02			WDS 2.5%
	Total # of fires	543,716					

Table E.6: Pyrome 5. R5 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	790.48	AVG PDSI	AVG GDP [US dollars]	2,013.19 ± 1,628.09
	AVG # of Fires	325,829.5	AVG Water deficit [mm]	AVG Population density [ppp/km ²]	27.13 ± 75.47
	AVG Size	7.79	AVG Temperature [K°]	AVG GDP [US dollars]	2,180.17 ± 773.76
	AVG Perimeter	12.38	AVG Max temperature [K°]	AVG Population density [ppp/km ²]	23.73 ± 32.46
			AVG Precipitation [in]	AVG Accessability [min]	283.09 ± 196.99
R6			AVG PDSI	AVG GDP [US dollars]	1,391.22 ± 597.51
	AVG Duration	5.39	AVG Water deficit [mm]	AVG Population density [ppp/km ²]	17.27 ± 32.07
	AVG Expansion	0.95	AVG Temperature [K°]	AVG Accessability [min]	372.31 ± 266.18
	AVG Perimeter/Area	1.70	AVG Max temperature [K°]	AVG GDP [US dollars]	1,880.84 ± 762.73
			AVG Precipitation [in]	AVG Population density [ppp/km ²]	51.8 ± 86.96
	N° of cells (res 1°)	188	AVG PDSI	AVG Accessability [min]	118.02 ± 66.93
	Total # of fires	5,213,272	AVG Water deficit [mm]	AVG GDP [US dollars]	1,878.83 ± 618.68
			AVG Temperature [K°]	AVG Population density [ppp/km ²]	21.87 ± 29.55
			AVG Max temperature [K°]	AVG Accessability [min]	262.77 ± 177.33
			AVG Precipitation [in]		

Table E.7: Pyrome 6. R6 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	3.50	AVG PDSI	AVG GDP [US dollars]	39,666.61 ± 10,112.52
	AVG # of Fires	961.2	AVG Water deficit [mm]	AVG Population density [ppp/km ²]	0.17 ± 2.10
	AVG Size	33.94	AVG Temperature [K°]	AVG GDP [US dollars]	22,476.51 ± 3,260.94
	AVG Perimeter	33.61	AVG Max temperature [K°]	AVG Population density [ppp/km ²]	0.07 ± 0.17
	AVG Duration	12.81	AVG Precipitation [in]	AVG Accessability [min]	3,260.62 ± 1761.49
R7			AVG PDSI	AVG GDP [US dollars]	39,666.61 ± 10,112.52
	AVG Expansion	1.54	AVG Water deficit [mm]	AVG Population density [ppp/km ²]	0.17 ± 2.10
	AVG Perimeter/Area	2.10	AVG Temperature [K°]	AVG GDP [US dollars]	22,476.51 ± 3,260.94
	N° of cells (res 1°)	358	AVG Max temperature [K°]	AVG Population density [ppp/km ²]	0.07 ± 0.17
	Total # of fires	13,500	AVG Precipitation [in]	AVG Accessability [min]	3,260.62 ± 1761.49

Table E.8: Pyrome 7. R7 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features				Demographic features				Land Cover %
	AVG Frequency	9.44	AVG PDSI	41.15 ± 120.84	AVG GDP [US dollars]	41,532.5 ± 5,512.08	CHO	30.6%			
	AVG # of Fires	15,612	AVG Water deficit [mm]	296.31 ± 206.14	AVG Population density [pop./km ²]	22.48 ± 80.22	GRS	24.7%			
	AVG Size	1.68	AVG Temperature [K]	281.53 ± 9.67	AVG Accessibility [min]	213.42 ± 313.19	WDS	18.4%			
	AVG Perimeter	5.36	AVG Max temperature [K°]	295.74 ± 8.45	AVG GDP [US dollars]	4,227.29 ± 2,403.82	ENC	7.4%			
			AVG Precipitation [in]	0.06 ± 0.02	AVG Population density [pop./km ²]	218.65 ± 419.32	DBF	6.6%			
			AVG PDSI	-115.7 ± 111.05	AVG GDP [US dollars]	17,577.48 ± 6,071.45	WDS	28.7%			
			AVG Water deficit [mm]	134.54 ± 81.68	AVG Population density [pop./km ²]	247.83 ± 294.98	SAV	23.7%			
			AVG Temperature [K°]	288.20 ± 6.3	AVG Accessibility [min]	175.77.48 ± 6,071.45	GRS	13.1%			
			AVG Max temperature [K°]	298.21 ± 5.11	AVG GDP [US dollars]	17,577.48 ± 6,071.45	ENC	10.5%			
			AVG Precipitation [in]	0.19 ± 0.09	AVG Population density [pop./km ²]	221.36 ± 161.5	CHO	6.6%			
			AVG PDSI	-27.41 ± 109.16	AVG GDP [US dollars]	5,275.18 ± 1,795.85	SAV	5.1%			
			AVG Water deficit [mm]	68.16 ± 148.52	AVG Population density [pop./km ²]	129.87 ± 463.14	GRS	64.2%			
			AVG Temperature [K°]	276.99 ± 9.98	AVG Accessibility [min]	222.1 ± 229.9	CHO	19.9%			
			AVG Max temperature [K°]	288.22 ± 10.63	AVG GDP [US dollars]	30,478.88 ± 10,477.64	DBF	9.2%			
			AVG Precipitation [in]	0.06 ± 0.02	AVG Population density [pop./km ²]	174.59 ± 270	ENC	6.5%			
			AVG PDSI	-155.53 ± 188.71	AVG GDP [US dollars]	33.94 ± 22.04	GRS	6.2%			
			AVG Water deficit [mm]	442.91 ± 426.27	AVG Population density [pop./km ²]		CHO	45.8%			
			AVG Temperature [K°]	278.53 ± 13.21	AVG Accessibility [min]		MRS	17.6%			
			AVG Max temperature [K°]	292.92 ± 12.21	AVG GDP [US dollars]		SAV	13.3%			
			AVG Precipitation [in]	0.04 ± 0.04	AVG Population density [pop./km ²]		GRS	6.5%			
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]		ENC	6.2%			
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						
			AVG PDSI	-141.85 ± 153.42	AVG GDP [US dollars]						
			AVG Water deficit [mm]	138.76 ± 194.18	AVG Population density [pop./km ²]						
			AVG Temperature [K°]	282.98 ± 6.83	AVG Accessibility [min]						
			AVG Max temperature [K°]	293.85 ± 8.09	AVG GDP [US dollars]						
			AVG Precipitation [in]	0.08 ± 0.03	AVG Population density [pop./km ²]						

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	77.37	AVG PDSI	AVG GDP [US dollars]	43,651.49 ± 4,498.78
	AVG # of Fires	24,574.9	AVG Water deficit [mm]	AVG Population density [ppt]/km2	0.09 ± 12.63
	AVG Size	106.54	AVG Temperature [K]	AVG Accessibility [min]	1,252.88 ± 635.63
			AVG Max temperature [K°]		
			AVG Precipitation [in]		
R11	AVG Perimeter	45.83	AVG PDSI	AVG GDP [US dollars]	8,835.03 ± 2,773.3
	AVG Duration	5.56	AVG Water deficit [mm]	AVG Population density [ppt]/km2	2.19 ± 6
	AVG Expansion	7.11	AVG Temperature [K]	AVG Accessibility [min]	424.13 ± 235.29
			AVG Max temperature [K°]		
			AVG Precipitation [in]		
	AVG Perimeter/Area	0.98	AVG PDSI	AVG GDP [US dollars]	13,733.36 ± 3,175.58
	N° of cells (res 1°)	294	AVG Water deficit [mm]	AVG Population density [ppt]/km2	3.41 ± 11.7
	Total # of fires	368,624	AVG Temperature [K]	AVG Accessibility [min]	271.37 ± 166.42
			AVG Max temperature [K°]		
			AVG Precipitation [in]		

Table E.12: Pyrome 11. R11 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	1175.74	AVG PDSI	AVG GDP [US dollars]	2,770.2 ± 2,727.32
	AVG # of Fires	250394.75	AVG Water deficit [mm]	AVG Population density [ppt]/km2	25.27 ± 127.8
	AVG Size	2.94	AVG Temperature [K]	AVG Accessibility [min]	297.41 ± 231.97
			AVG Max temperature [K°]		
			AVG Precipitation [in]		
R12	AVG Perimeter	7.61	AVG PDSI	AVG GDP [US dollars]	1,425.04 ± 543.39
	AVG Duration	4.34	AVG Water deficit [mm]	AVG Population density [ppt]/km2	35.17 ± 165.24
	AVG Expansion	0.55	AVG Temperature [K]	AVG Accessibility [min]	160.91 ± 96.86
			AVG Max temperature [K°]		
			AVG Precipitation [in]		
	AVG Perimeter/Area	2.82	AVG PDSI	AVG GDP [US dollars]	786.45 ± 102.88
	N° of cells (res 1°)	179	AVG Water deficit [mm]	AVG Population density [ppt]/km2	22.14 ± 75.8
	Total # of fires	4,006,316	AVG Temperature [K]	AVG Accessibility [min]	343.77 ± 145.15
			AVG Max temperature [K°]		
			AVG Precipitation [in]		

Table E.13: Pyrome 12. R12 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	307.14	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	SAV 41.5% GRS 27.1% CRO 10.7% WDS 9.4% MFS 7.5%
	AVG # of Fires	132,845	-92.44 ± 202.42 296.33 ± 287.1 297.33 ± 2.54 305.89 ± 2.23 0.16 ± 0.13	AVG GDP [US dollars] AVG Population density [pp]/km2	EBP 34.8% CRO 24.4% WDS 16.3% SAV 11.9% MFS 5%
	AVG Size	1.83	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	EBP 34.8% CRO 24.4% WDS 16.3% SAV 11.9% MFS 5%
	AVG Perimeter	5.81	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	EBP 34.8% CRO 24.4% WDS 16.3% SAV 11.9% MFS 5%
R13	AVG Duration	4.15	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	CRO 78.4% GRS 9.5% SAV 2.6% WBS 2.6% MFS 2.5%
	AVG Expansion	0.38	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	CRO 78.4% GRS 9.5% SAV 2.6% WBS 2.6% MFS 2.5%
	AVG Perimeter/Area	3.66	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	CRO 78.4% GRS 9.5% SAV 2.6% WBS 2.6% MFS 2.5%
	N° of cells (res 1°)	503	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	EBP 38.8% SAV 22.3% NV 13.1% WBS 9.4% GRS 8%
	Total # of fires	2,125,520	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [in]	AVG GDP [US dollars] AVG Population density [pp]/km2	EBP 38.8% SAV 22.3% NV 13.1% WBS 9.4% GRS 8%

Table E.14: Pyrome 13. R13 pyrome and regimes description.

Group	Fire characterization	Sub-groups	Climate features	Demographic features	Land Cover %
	AVG Frequency	6.25	AVG PDSI AVG Water deficit [mm] AVG Temperature [K] AVG Max temperature [K°] AVG Precipitation [m]	AVG GDP [US dollars] AVG Population density [ppl./km2]	WDS 36.7% SAV 27% MFS 13.7% ENC 4% GRS 3.7%
	AVG # of Fires	6,090.4	R14-a Area 4,603,000 km ²	1.03 ± 11.34	
	AVG Size	8.98	AVG PDSI AVG Water deficit [mm] AVG Temperature [K°] AVG Max temperature [K°] AVG Precipitation [m]	AVG Accessability [min] AVG GDP [US dollars]	2,135.49 ± 1,716.49
	AVG Perimeter	14.98	R14-b Area 1,551,000 km ²	AVG Population density [ppl./km2]	SAV 36.8% WDS 34.5% OSL 10.9% ENC 9.8% WBS 3.7%
R14	AVG Duration	8.32	AVG PDSI AVG Water deficit [mm] AVG Temperature [K°] AVG Max temperature [K°] AVG Precipitation [m]	AVG Accessability [min] AVG GDP [US dollars]	1,973.91 ± 1,200.63
	AVG Expansion	0.66	R14-c Area 344,000 km ²	AVG Population density [ppl./km2]	28,543.96 ± 2,858
	AVG Perimeter/Area	3.03	AVG PDSI AVG Water deficit [mm] AVG Temperature [K°] AVG Max temperature [K°] AVG Precipitation [m]	AVG Accessability [min]	SAV 50% OSL 45.2% GRS 4.3%
	N° of cells (res 1°)	967	R14-d Area 192,000 km ²	AVG GDP [US dollars] AVG Population density [ppl./km2]	87,872.49 ± 26,410.96
Total # of fires		91,356	AVG PDSI AVG Water deficit [mm] AVG Temperature [K°] AVG Max temperature [K°] AVG Precipitation [m]	AVG Accessability [min] AVG GDP [US dollars]	WDS 43.7% ENC 24.1% MFS 12.1% SAV 10.2% WBS 7.5%
			R14-e Area 186,000 km ²	AVG Population density [ppl./km2]	SAV 69.6% WDS 21.6% PWL 5% WBS 3.3%

Table E.15: Pyrome 14. R14 pyrome and regimes description.

E.2 Data & Code availability.

All data products and sub-products were recorded and organized into a series of GeoTif and shapefiles, ready to share with other researchers and the community in an open-source repository <http://www.github.com/cpaismz89/Vulcano> and as part of Google Earth Engine assets. The code to reproduce every step of our research is available as a series of Python scripts and Jupyter notebooks for visualization convenience.