Landscape Optimization for Prescribed Burns in Wildfire Mitigation Planning

Weizhe Chen
University of Southern California
Los Angeles, United States
weizhech@usc.edu

Eshwar Prasad Sivaramakrishnan
University of Southern California
Los Angeles, United States
esivaram@usc.edu@usc.edu

Bistra Dilkina
University of Southern California
Los Angeles, United States
dilkina@usc.edu

ABSTRACT

Wildfires have increased in extent and severity, and are posing a growing threat to people's well-being and the environment. Prescribed burns (burning on purpose parts of the landscape) are one of the key mitigation strategies available to reduce the potential damage of wildfires. However, where to conduct prescribed burns has long been a problem for domain experts. With the advancement of forest science, weather science, and computational modeling, there produced powerful fire simulators that can help inform how wildfires will start and grow. In this paper, we model the problem of selecting where to perform a set of prescribed burns across a large landscape into a multi-objective optimization problem. We build a surrogate objective function from simulation data and solve the multi-objective optimization problem with genetic algorithms. We name our solution as Spatial Multi-Objective for Prescribed Burn (SMO-PB). We also investigate three variants of the approach that further consider spatial fairness. With a case study of Dogrib, Canada, we show that our formulations can successfully provide solutions capable of real world deployment, and showed how fairness can be reached without diminishing the performance a lot.

CCS CONCEPTS

• **Social and professional topics** → **Sustainability**; • **Applied computing** → **Multi-criterion optimization and decision-making**.

KEYWORDS

Prescribed Burn, Wildfire, Genetic Algorithm, Multi-objective Optimization, Fairness, Landscape Optimization

ACM Reference Format:
Weizhe Chen, Eshwar Prasad Sivaramakrishnan, and Bistra Dilkina. 2022. Landscape Optimization for Prescribed Burns in Wildfire Mitigation Planning. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS) (COMPASS '22), June 29-July 1, 2022, Seattle, WA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3530190.3534816

COMPASS '22, June 29-July 1, 2022, Seattle, WA, USA

ACM ISBN 978-1-4503-9347-8/22/06.
https://doi.org/10.1145/3530190.3534816

1 INTRODUCTION

Although technology has greatly improved in recent years, wildfire is a consistent threat to the environment and the quality of people's lives. In the last two years, in California alone, there were about 7 million acres burned because of wildfire, with tens of thousands of infrastructures destroyed and dozens of lives lost [23]. Wildfires can be caused by lightning, volcanic activity, spark, or human carelessness, so it is really hard to stop wildfire from happening. Even worse, research also indicates that the probability of these fires has increased in recent years due to the effects of climate change on temperature, precipitation levels, and soil moisture [27, 32].

Taking actions to mitigate the wildfire risk ahead of time is one of the key tasks for fire and forest agencies to do. Specifically, fuel management is a commonly accepted way to reduce the intensity and severity of wildfires. There are multiple targeted fuel treatment activities. For example, prescribed burns, thinning, and mechanical treatments are popular and accepted by both researchers and domain experts. In this work, we mainly consider prescribed burn as the way to mitigate wildfire risks.

A prescribed burn is a planned burn by the fire department as a controlled application of fire in order to greatly reduce fire hazards. While prescribed burning is one of the most important tools used to manage fire today, there are still several factors, such as the location of the prescribed burn, that can affect the effectiveness of a prescribed burn plan. Landscape-level siting for prescribed burns is traditionally planned by domain experts such as site managers and local fire departments with the main goal of reducing future wildfire impacts (e.g., area burned, CO2 produced, infrastructure damaged). The total number of possible prescribed burn configurations can be exponentially large with respect to the number of candidate locations. Although the prescribed burn plan does not need to be planned quickly, and only needs to be planned once every several months, the large decision space combined with a complex combination of objectives, makes prescribed burn landscape optimization a tough task, as the planner needs to take into account various metrics at the same time. Currently, prescribed burns are normally planned by giving an independent score for each unit and choosing the best few units as a solution. The score itself is very often set by human experts based on domain knowledge, rather than data-driven. An additional challenge is posed by the fact that the simulation for wildfire is relatively slow, and can take up to minutes to run a group of simulations that is representative enough for one specific prescribed burn plan. So it is impossible to use simulators during the optimization process, which makes enumeration over solutions for an optimal plan impossible.

In this paper, we propose a data-driven simulation-optimization approach for the landscape optimization for prescribed burns, aiming to better consider the different objectives simultaneously. For the simulation data for wildfires without any prescribed burns, we first build a spread-tree-based surrogate function to quickly estimate the benefit we can get from a specific prescribed burn plan without running the simulation, and use mathematical programming to formulate the problem into a multi-objective knapsack problem. We name our solution as Spatial Multi-Objective for Prescribed Burn (SMO-PB). To make our solution more acceptable to stakeholders and residents of the affected areas, we integrate considerations of how the benefits (i.e., reduced wildfire hazards) are spatially distributed. We give three formulation variants that introduce the potential fairness concerns for field deployment. The workflow is shown in Figure 1. We use a case study of Dogrib, a 79,611 ha large surface in the Rocky Mountains, Canada to show that our algorithm can generate the Pareto frontier efficiently, and show how the fairness-aware formulations do not compromise overall wildfire hazard reduction. We also use an ablation study to give some insights on how to better the formulation results with more, though possibly inexact, given candidates.

Our main contributions are: 1) We formulate the landscape optimization for prescribed burn planning into a four-stage problem. Previous work normally disregarded one or two of these stages. 2) We provide a general group of multi-objective optimization formulations for landscape optimization. 3) We introduce the fairness concerns into the prescribed burn problem. 4) We use a case study of Dogrib to show that our workflow can generally work well both from the computational and the solution quality perspective.

2 RELATED WORK

Wildfire has long been a focus for fire and forest researchers, and there are several wildfire simulators built to make the decision-making more accurate and easier. For example, FARSITE [12] was one of the most used simulators, because of its powerful simulation and long-term collaboration with fire researchers. QUICFIRE [20] was created in recent years to model 3D time-resolved fire behavior. Cell2Fire [24] is created for modeling the exact spread tree of the wildfire in a calculation-efficient way.

In this paper, we are focusing on landscape-level optimization for prescribed burns in wildfire mitigation, where a limited subset of prescribed burn areas is selected from a larger set of candidates to optimize their joint complementary effect on reducing the wildfire hazards across the landscape. There were a few works specifically on planning prescribed burns. Hanselka et al. [15] summarized the principles for designing a prescribed burn plan for domain experts' management. Cowell et al. [6] proposed a score-based ranking system for prescribed burn planning. Kim et al.[17] evaluated a set of different fuel management settings in certain patterns (i.e., dispersed, clumped, regular, and random patterns), and with a case study in Oregon, they concluded that the proposed method can only marginally alter the size and severity of future wildfires under tight budgetary conditions. Others [5, 13] instead studied the benefit of placing a series of parallel strips-like fuel treatments, and letting the fire propagate perpendicularly to their placement. Russo et al.[28] studied the placement of firebreaks to control fire spread. Pais et al. [25] proposed a useful surrogate function for potential savings derived from wildfire simulation data and advocated its general usage in optimization for fuel treatment. We use the surrogate function of [25] as the basis for the optimization objectives in this paper, but for the optimization part, Pais and colleagues only considered a connected area for protection, while we are considering general combinations of candidate prescribed burn areas, and only a single objective, while we are considering multiple objectives.

Matsypura et al. [22] studied prescribed burn placement by network-based optimization on an abstraction graph, and used a case study of Hawkesbury to validate their solution quality. However, the scalability of their algorithm is only assessed on graphs with hundreds of nodes, which is far from the real-world need. The work by Alcasena et al. [2] is probably the closest research paper to this paper, in which they also formulate the problem into a multi-objective optimization. Their solution is a weight-sampling-based method to give a production possibility frontier (PPF), while our solution is to directly solve the multi-objective optimization. Furthermore, we also consider how fairness concerns in this problem can be solved.

Multi-objective optimization is a well studied topic in operations research and among optimization researchers because of its potential applications, e.g., dial a ride problem (DARP) [14] among many others. In this research, evolutionary multi-objective optimization (EMO) methodologies have amply shown their niche in finding a set of well-converged and diversified non-dominated solutions since the beginning of the 1990s. Non-dominated Sorting Genetic Algorithm (NSGA)-II [9] is one of the most popular algorithms used in this domain, because of the good solution quality and the early time it is proposed. Later on, researchers improved the NSGA-II to NSGA-III, making it one of the state-of-the-art approaches [8]. Several extensions for NSGA-II are proposed, and many of them are also applicable to NSGA-III, resulting in a group of variants of NSGA-III (e.g., U-NSGA-III [31], R-NSGA-III [29]). Other lines of work in multi-objective optimization mostly include variants of the multi-objective evolutionary algorithm (MOEA): MOEA/D[34] and AGE-MOEA[26], and some other algorithms that do not have as many variants like SPEA2 (Strength Pareto Evolutionary Algorithm-2) [36] and Two-Archive Evolutionary Algorithm for Constrained Multiobjective Optimization (C-TAEA) [19]. The initial formulation in this paper is a 0-1 multi-objective knapsack problem [21]. Besides the above approaches for the general problem, past works [4, 10, 18, 30] have proposed a few exact solution methods that can be efficient in some special cases. Recently, Zhou et al. [35] also proposed the use of reinforcement learning for multi-objective optimization. In this paper, we use NSGA-III, U-NSGA-III, C-TAEA because of their wide applicability and good public availability.

In this paper, we also introduce fairness to the landscape optimization problem for prescribed burn planning. Fairness considers giving similar outcomes to different groups, as defined by some relevant features. This is an important concept studied for a long time, but recently there have been considerable efforts to incorporate it into computational models. [33] started to address fairness by quantifying the variance of the result methods get. [16] summarized some of the properties a fairness metric usually needs to satisfy, and then proposed a new division-based index for fairness. [1, 11] also studied fair resource allocations in infrastructure in bike-sharing

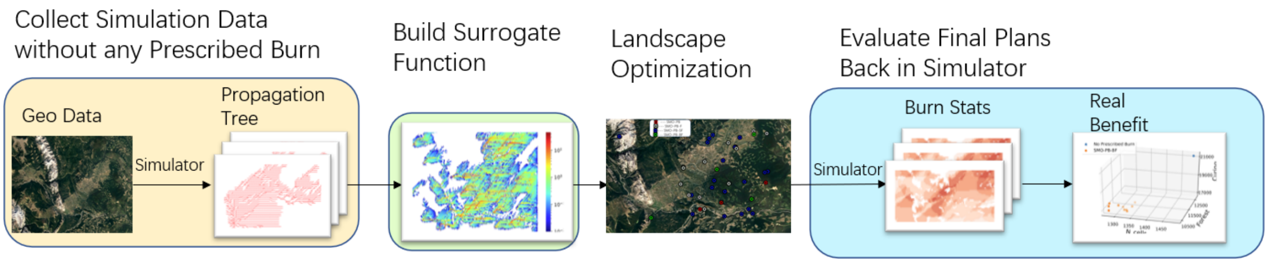

Figure 1: The workflow of our landscape optimization solution for prescribed burn in wildfire mitigation planning.

Figure 2: The satellite image of Dogrib, Canada.

and community evacuation planning and gave different solutions for their application domain. [3] studied fair resource allocation, used a welfare-based dominance constraint, and used a case study of workload allocation to show that such a problem can be solved efficiently. The fairness concern in our problem can also be seen as a fair resource allocation problem, however, we used a different definition and solution because our goal is different.

3 PRELIMINARY

3.1 Landscape Optimization for Prescribed Burns

Prescribed burning is a commonly used technique to help reduce the potential threat of wildfires. Unlike real-time firefighting, prescribed burns are done before any wildfire starts, and are under the control of a local fire department. It is useful for clearing downed trees, preventing tree diseases, and most importantly, reducing future wildfire hazards. Because this is done after careful planning and condition controls, the fire experts can accurately control and predict in which specific area they can do a prescribed burn. However, because of the large resource requirement for every prescribed burn with careful condition control, local fire managers cannot do many prescribed burns, and choosing from the large candidate set of where exactly they should do prescribed burns is a key challenge for them. At the same time, prescribed burns across a landscape have complementary and compounding effects on the wildfire hazard, and hence should be planned jointly as a configuration of burns instead of planned independently.

Specifically, for a landscape area A, we are given a list of objective functions $F = \{f_1, f_2, f_3, \cdots, f_M\}$ for our wildfire protection, where M is the number of objectives. Each of these objectives is a function that takes as input a prescribed burn configuration x, and $f_m(x)$ represents some real-world metric for wildfire hazard reduction due to applying plan x that is well accepted in the wildfire domain, for example, expected reduction in the infrastructure destroyed, the CO2 emissions, or the damage to local ecosystems by wildfires. In addition to measuring the overall wildfire hazard reduction $f_m(x)$ in the landscape A after the prescribed burn plan x is applied, we also have access to $f_m(a, x)$ which specifies the same wildfire hazard metric under plan x but for a specific location/subarea $a \in A$. We also have a set of K different candidate prescribed burn locations $X = \{X_1, X_2, X_3, \cdots, X_K\}$ given by domain experts considering the real deployment needs like the distance to local fire departments, the slope, etc. Our task is to determine which subset $x \subset X$ one should choose so that the objectives $F(x)$ we get for this prescribed burn configuration are optimized. For notation simplicity, x is encoded in a Boolean vector form such that $x_k = 1$ means the k-th element in X is selected for the plan, and 0 otherwise. Because of the limited resources in the real world, we are also given L different groups of weights $W_{K\times L}$ corresponding to a group of limited total budgets $B = (b_1, b_2, \cdots, b_L)$ for prescribed burns. The limited resources can be the number of burns, the money for prescribed burns, the time effort for prescribed burns, etc. This is a constrained multi-objective optimization, and we want to calculate the Pareto frontier solutions with respect to the objectives.

Note that we require the prescribed burn candidates to be given by experts because in real world deployment, the size of each prescribed burn should be limited. Since we do not limit the number of candidates here, we assume the candidates can contain whatever the domain experts think can be a reasonable unit in a specific plan, and do not consider the case that the given candidate set is incomplete.

3.2 Surrogate Objective Function

The most accurate way to compute the wildfire hazards in a landscape and their reduction under a proposed prescribed burns plan is to use a wildfire simulator. Wildfire simulations are computed by treating the landscape as a regular grid of cells, sampling weather conditions and ignition points, and simulating the fire spread from cell to cell. For a given simulation s, we obtain the wildfire impact or damage on each cell i for each objective metric m, denoted by

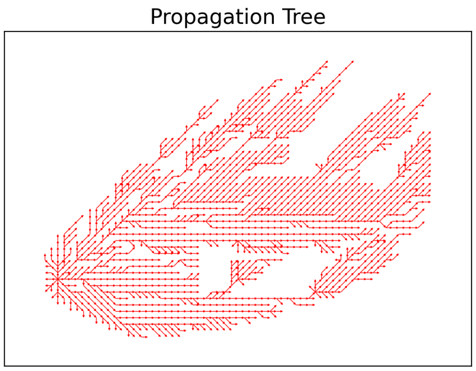

(a) An example of fire propagation tree generated by the simulator.

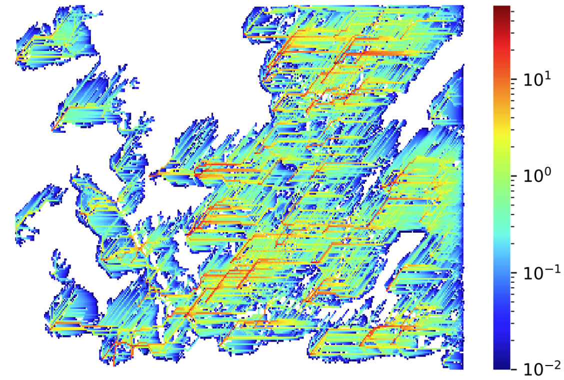

(b) The DPV heatmap for unit node value of the Dogrib area generated from simulation data used in this paper.

Figure 3: We get the propagation tree data like figure (a) from the simulator, to calculate the surrogate function for further optimization usage, shown in figure (b). Note that (b) is not the result directly generated from (a), but the average result of 100 propagation trees like (a).

$NV_m^s(i)$. A representative set of simulations give an estimate of the expected wildfire impacts.

During the optimization, the objective functions would take a long time for wildfire simulators to run before the final result can be given. For example, the Cell2Fire simulator takes at least one second to run one simulated wildfire, which is unacceptable for any optimization algorithm since most of them will run this objective function calculation tens of thousands of times during the optimization. Thus, computing the objectives under a proposed plan through simulation is not reasonable during our optimization. In this paper, we use a modified version of *downstream protection value* [25] to calculate a surrogate function.

Downstream protection value (DPV) [25] is a recently proposed metric that measures and ranks the impact of treating a unit cell of the landscape by estimating what is burned downstream from it, starting from a burning cell, in a representative sample of wildfires. During the simulation, the fire spread process is recorded between the cells and results in a directed graph of burning cells with nodes representing cells, as illustrated in Figure 3a. Then, a directed shortest-path-tree is generated based on the directed graph. The DPV value of a cell/node i captures the prevented wildfire impacts if the cell i was not allowed to burn or spread wildfires (e.g. due to a prescribed burn), and is calculated as the average value over all simulations of the sum of all the node values (NV) of all the children of the cell i on the shortest path tree in each wildfire simulation s:

$$DPV_m(i) = \frac{1}{n_{sims}} \sum_{s=1}^{n_{sims}} \sum_{j \in Subtree_s(i)} NV_m^s(j) \tag{1}$$

where $NV_m^s(i)$ is the node value of node i under the m-th objective in the s-th simulation, $Subtree_s(i)$ is the set of cells/nodes in the fire-spread subtree of node i in the s-th simulation, and n_{sims} denotes the total number of simulations.

With the DPV approximation, we can think of the objective functions for each prescribed burn plan x as:

$$\forall m \in \{1, 2, \cdots, M\}, f_m(x) = \sum_{k=1}^{K} DPV_m(k) * x_k \tag{2}$$

to overlaps of the mitigation benefits of individual prescribed burns x_k. The precise encoding is presented in the next section.

This surrogate function relies on an assumption that the wildfire in the simulation can only spread through exactly what is on the provided fire spread tree, but not any other edge. This is clearly not true in the real world and can lead to a gap between the estimated wildfire hazard reduction value and the actual benefit, but still, this is so far one of the most accurate surrogate functions that can be computed efficiently. Generally, using surrogate functions requires a re-run on the simulator after the whole optimization process to get the real benefit.

4 MATHEMATICAL MODEL

4.1 Spatial Multi-Objective Optimization for Prescribed Burns

We start with giving our basic multi-objective optimization formulation:

$$\begin{aligned} &\max_x F(x) = (f_1(x), f_2(x), f_3(x), \cdots, f_M(x)) \\ &s.t. Wx \leq B \\ &\forall k, x_k \in \{0, 1\} \end{aligned} \tag{3}$$

where we aim to maximize the M metrics of wildfire impact reductions, subject to L resource constraints, under binary decision variables for each prescribed burn candidate x_k.

For a special case with $M = 1$, the problem turns into a standard integer programming, or more specifically, a knapsack problem. One possible solution is to use dynamic programming. However, the knapsack problem is also NP-hard, and cannot be efficiently solved in real world-scale large problems.

In the general case, our problem is naturally a multi-objective combinatorial optimization problem, also known as the multi-objective subset selection problem, and more specifically, a *0-1 multi-objective knapsack problem* [21]. We call this formulation *Spatial Multi-Objective Optimization for Prescribed Burns* (SMO-PB).

Next, we provide the exact mathematical formulation that captures the DPV objectives in wildfire mitigation:

$$
\begin{aligned}
\max_{x,z} f_1(z) &= \frac{1}{n_{sims}} \sum_{i=1}^{|A|} \sum_{s=1}^{n_{sim}} NV_1^s(i) * z_{i,s} \\
f_2(z) &= \frac{1}{n_{sims}} \sum_{i=1}^{|A|} \sum_{s=1}^{n_{sim}} NV_2^s(i) * z_{i,s} \\
&\cdots \\
f_M(z) &= \frac{1}{n_{sims}} \sum_{i=1}^{|A|} \sum_{s=1}^{n_{sim}} NV_M^s(i) * z_{i,s} \\
s.t. Wx &\le B \\
\forall i, s, z_{i,s} &\le \sum_{k \in Parent_s(i) \cap X} x_k \\
\forall i, s, z_{i,s} &\in [0,1] \\
\forall k, x_k &\in \{0,1\}
\end{aligned}
\tag{4}
$$

where $z_{i,s}$ is a group of auxiliary variables that encodes whether a cell i will be burned in simulation s after applying the corresponding prescribed burn plan x, and $Parent_s(i)$ gives a set of the parents of cell/node i on the spread tree generated by simulation s.

The biggest limitation of this formulation is that we need a simulator that can provide the propagation tree for the simulations. Even when the fire spread information is not available, usually any wildfire simulator will at the minimum provide for each simulation the cell of the ignition point and fire impacts in each cell captured by the node values $NV_m^s(i)$. In that case, our formulation still applies but $Parent_s(i)$ is only the cell of the ignition point of fire simulation s.

4.2 Fair Spatial Multi-Objective Optimization for Prescribed Burns

Fairness considers giving similar outcomes to different groups, as defined by some relevant features. In the context of our application in wildfire hazard mitigation, we found that it is quite common to have a plan in which selected burns are very close to each other, because typically the high threat areas are in a cluster. This makes the final benefit each cell received from the prescribed burn not evenly distributed. In reality, other parts of the park also have to suffer the smoke and other potential damage during the prescribed burn without receiving benefits. From the perspective of fire managers, they want to make most people get a benefit by making the benefit distributed more evenly, and hopefully more residents in the potential prescribed burn area are more willing to support the prescribed burns.

We formulate this fairness concern into the optimization itself, such that in addition to the overall benefits (total wildfire impacts' reduction across the landscape) we also measure the distribution of these mitigation benefits to each cell / subarea of the landscape. For clarity, we will call the current objectives we get directly from the simulations the "main objectives". We proposed three extended formulations that try to find a solution that does not reduce much the main objectives, while also providing a more fair solution in terms of the distribution of the mitigation benefits.

4.2.1 Spatial Area Fairness. Given cell i and plan x, we define $f_m(i,x)$ as the reduction in wildfire impact m in cell i under plan x. Based on the variables in the formulation in Eq.4, one can compute $f_m(i,x)$ as $\frac{1}{n_{sims}} \sum_{s=1}^{n_{sim}} NV_m^s(i) * z_{i,s}$ using the auxiliary variables z.

We define the fairness metrics in terms of the lowest quantile of the benefits to individual cells under the prescribed burn plan. For a length l sequence, the lowest quantile is the value v such that exactly $\frac{l}{4}$ elements of the sequence have a value less than v. In this way, instead of optimizing for the worst-off cell which usually will have 0 benefits for a plan x, we focus look at the 25%-th cell:

$$
e_m(x) = \frac{1}{z_m} lowest\_quantile_{i \in A} f_m(i,x) \tag{5}
$$

where e_m is the fairness metric corresponding to main objective m, and z_m is a normalizing parameter to make the fairness metric value always between 0 and 1. We pre-compute z_m before the multi-objective optimization, by optimizing a single objective without considering other objectives:

$$
\begin{aligned}
z_m \triangleq \max_x & f_m(x) \\
s.t. & Wx \le B \\
& \forall k, x_k \in \{0,1\}
\end{aligned}
\tag{6}
$$

Hence, we have M main objectives, now we have an additional M fairness objectives $E = \{e_1, e_2, \cdots, e_M\}$ with $Z = (z_1, z_2, \cdots, z_M)$.

4.2.2 Spatial Multi-Objective Optimization for Prescribed Burns with Fairness (SMO-PB-F). We first directly introduce the fairness functions as another group of objective functions with the original objective functions into formulation Eq. 3:

$$
\begin{aligned}
\max_x & F(x) = (f_1(x), f_2(x), f_3(x), \cdots, f_M(x)) \\
& \boldsymbol{E(x) = (e_1(x), e_2(x), \cdots, e_M(x))} \\
s.t. & Wx \le B \\
& \forall k, x_k \in \{0,1\}
\end{aligned}
\tag{7}
$$

This formulation results in a solution with *(2M)-Dimensional* Pareto frontier. Domain experts and site managers can choose one of the Pareto-front solutions to deploy. We call this formulation Spatial Multi-Objective Optimization for Prescribed Burns with Fairness (SMO-PB-F).

4.2.3 Spatial Multi-Objective Optimization for Prescribed Burns with Sum Fairness (SMO-PB-SF). Above, SMO-PB-F is a formulation that keeps the greatest flexibility in selecting a solution. However, for current multi-objective optimization algorithms, the scalability is negatively impacted by the growth in the number of objectives. In addition to being slow, since the number of objective functions is doubled it will also require a lot of human post-processing in selecting a plan from the Pareto front. Instead of adding M fairness objectives, we can use the sum for the fairness metrics and add exactly one objective function into Eq. 3:

$$\begin{aligned} &\max_x F(x) = (f_1(x), f_2(x), f_3(x), \cdots, f_M(x)) \\ &\qquad \boldsymbol{e_1(x) + e_2(x) + \cdots + e_M(x)} \\ &s.t. Wx \leq B \\ &\qquad \forall k, x_k \in \{0, 1\} \end{aligned} \tag{8}$$

This formulation can help in the context of scalability for the formulation and can give a solution that considers fairness in all objectives given that $e_m(x)$ are normalized. Also, this formulation can be easily extended to a weighted sum version instead of uniform sum with the help of domain experts. We call this formulation Spatial Multi-Objective Optimization for Prescribed Burns with Sum Fairness (SMO-PB-SF).

4.2.4 Spatial Multi-Objective Optimization for Prescribed Burns with Bounded Fairness (SMO-PB-BF). Often in research related to fairness, one adds fairness constraints instead of adding objectives because most problems are solved as a single-objective problem and avoid considering a Pareto frontier. So instead of adding objectives to the Eq. 3 formulation, we alternatively add m constraints:

$$\begin{aligned} &\max_x F(x) = (f_1(x), f_2(x), f_3(x), \cdots, f_M(x)) \\ &s.t. Wx \leq B \\ &\qquad \boldsymbol{\forall m \in \{1, 2, \cdots, M\}, e_m(x) \geq \epsilon_m} \\ &\qquad \forall i, x_i \in \{0, 1\} \end{aligned} \tag{9}$$

where $\epsilon = (\epsilon_1, \epsilon_2, \epsilon_3, \cdots, \epsilon_M)$ is another hyperparameter class that defines a lower bound threshold on the quantile cell benefit. The larger the number, the more fair the solution will be.

This formulation is reasonable because it maintains exactly the same amount of objective functions as SMO-PB. This is beneficial in terms of runtime as well as visualization of the Pareto front. Furthermore, reducing the number of objectives naturally reduces the potential number of points on the Pareto fronts, and thus reduces the need for human experts to choose solutions among a large set. We call this formulation Spatial Multi-Objective Optimization for Prescribed Burns with Bounded Fairness (SMO-PB-BF).

Note that after adding the fairness concerns, the three variants we get are no longer a simple linear multi-objective optimization or a multi-objective knapsack problem, which makes the formulation much harder to solve; however, the current solver for the multi-objective optimization is good enough to solve these problems.

5 EXPERIMENT

5.1 Experiment Setting

We have tested our formulations in a real world-based dataset in Dogrib, Canada. We split the area into cells of $100m \times 100m$, and provide the features of these cells to the simulator for wildfire simulation. We used the Cell2Fire simulator [24] to simulate the fire 100 times with different ignition points over the whole area according to the average ignition probability of 10 different kinds of weather. The simulator is chosen because it is one of the very few simulators that provide detailed propagation graphs from the simulation, which enable the calculation of the DPV values. Specifically in this paper, we choose to optimize number of cells burned (N_cells), the forest burned (Forest), and the sum of the degree of curing of the cells burned (Curing). We also optimize their corresponding fairness metrics E, to optimize area fairness (E(N_cells)), forest fairness (E(Forest)), and curing fairness (E(Curing)). The forest burned per cell is an integer from 0 to 110, and the degree of curing per cell is an integer from 0 to 70, assigned by the wildfire simulator. [1] For candidate prescribed burn areas, we randomly select some cells within the cells that are burned in at least one of the simulations. For simplicity, we only consider the case that there is one budget limit and each candidate burn area has a uniform weight, i.e., $L = 1, \forall k \in \{1, 2, \cdots, K\}, w_{1,k} = 1$. We choose a budget of $B = 30$. Without further specificity, for SMO-PB-BF we set $\forall m, \epsilon_m = 5 \times 10^{-3}$.

Table 1: The runtime results of our formulations, with 500 candidates and 500 generations.

Runtimes(s)	SMO-PB	SMO-PB-F	SMO-PB-SF	SMO-PB-BF
NSGA-II	80.88	596.97	542.20	536.44
NSGA-III	90.93	664.24	657.62	655.63
U-NSGA-III	77.23	537.52	537.64	534.98
C-TAEA	87.65	712.86	673.76	650.14
PPF	17855.33	-	-	-

For the multi-objective optimization, we used the popular NSGA family algorithms: NSGA-II [9], NSGA-III [8], U-USGA-III [29], and C-TAEA [19]. We also compared our results with the PPF method used in [2]. For the reference direction needed in these genetic algorithms, Das-Dennis method [7] is used to sample 100 direction vectors. Note that all of these algorithms are sub-optimal algorithms, so we compare their results with similar runtime limits. We let all NSGA-group algorithms run for 500 generations, and C-TAEA for 1000 simulations since C-TAEA is twice as fast in every generation (as shown in Table 1). All experiments are tested on a Google Colab server with 13GB memory, and numbers are averaged from 3 different seeds.

5.2 Optimization Results

We show the convergence plot of different algorithms and different formulations in Fig. 4. We can see that 500 generations are enough for most scenarios to converge. As such, in later experiments without further specificity, the results are from 500 generations.

Table 1 shows the runtime of different algorithms and formulations. As expected, the fastest formulation is the SMO-PB, and we can see that other formulations are much slower than the original one. In Table. 2, we show the hypervolume results of different algorithms and different formulations. Although PPF performs slightly better in hypervolume than other algorithms in the SMO-PB formulation, it needs to enumerate different weight combinations between objectives taking several orders of magnitudes longer time, and is even intractable in our other formulations requiring more than 3 objectives. Among the other algorithms, we observe that for most of the settings different algorithms can achieve similar hypervolume

[1] The data was downloaded together with the open sourced simulator: https://github.com/cell2fire/Cell2Fire.

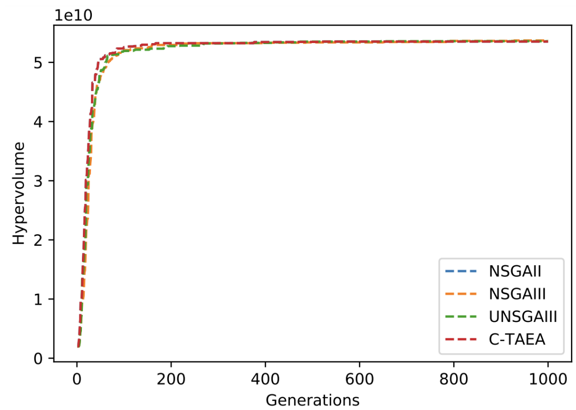

(a) Convergence for SMO-PB optimized by different algorithms.

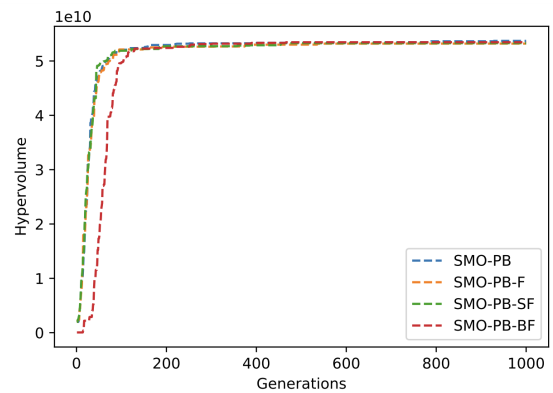

(b) Convergence plot for different formulations optimized by NSGA-III.

Figure 4: Convergence in hypervolume across different algorithms and different formulations. The hypervolume is calculated based on the three main objectives.

results for each of the formulations, and specifically, NSGA-III is the one that performs the best most of the time. So we chose to use NSGA-III to report further in-depth results.

In Table. 3, we show all of the results of NSGA-III with different formulations. We show the hypervolume with 3 different settings:

- 3 main objectives, namely number of cells burned
- 3 fairness objectives, i.e., $e1(x)$, $e2(x)$, $e3(x)$
- All 6 objectives with both the main objectives and the fairness objectives.

We can see that although the runtime for the SMO-PB formulation is much smaller than other formulations that have considered fairness, the hypervolume in terms of fairness objectives of the SMO-PB formulation is much smaller compared to other formulations. From this point of view, spending a reasonably longer runtime and finding a better solution that does not lose much on the original three objectives, and with much better fairness, is quite reachable. SMO-PB-BF gets the best fairness result when only considering the fairness objectives and the fewest number of points on the Pareto Fronts, because it has a strict constraint to ensure fairness, and loses a lot in the hypervolume of the 3 main objectives. So whether using the SMO-PB-BF model or using other fairness-concerned models should decide based on how many solutions a domain expert wants to choose from. In Fig. 5a, we show the 3-D plot of the corresponding Pareto fronts.

In Fig. 5b we show the 3-D projection of the Pareto frontier of the first two main objectives and the fairness objective for forest E(Forest), from different formulations (for one specific random seed). We can see that the whole Pareto frontier surface of the SMO-PB formulation is in the bottom half of all of the multi-objective space.

We randomly chose one plan from the Pareto front of each formulation. In Fig. 6, we show the benefit each cell gets from the plan as heatmaps. We see that although most solutions give benefit to the same area, the result from SMO-PB gives more extreme values while other formulations give more average values to everyone. Only SMO-PB-F and SMO-PB-BF have the biggest motivation to spend less in those better-rewarding areas, and protect those areas that are normally not covered (e.g., the middle part of SMO-PB-F, the lower left area of SMO-PB-BF). In Fig. 7, we show the actual plan in the corresponding area. While many candidates overlapped in different solutions, we can see that all of the candidates in the corners are only covered by the fairness-aware formulations. These distribution figures support our assumption that considering fairness can benefit more people.

Although SMO-PB-F takes the longest to run, it does not outperform the result of SMO-PB-SF, which means that simply using a sum for fairness objectives can be good in both runtime and solution quality. Therefore, we believe that for real world deployment where we have more objectives to consider than the two objectives considered in this section, SMO-PB-SF and SMO-PB-BF will be a better choice.

Figure 8 and Table 4 shows how our formulations perform with different numbers of candidate prescribed burn areas. We can see that with the increase in number, our solution quality also gets better. This is because our formulation does not generate other plans that are not in the candidates; therefore the more candidates, the more likely that the optimal plan is included in the candidate set. This means that for domain experts, it will be good if they can provide a large group of candidates even if they do not have such large budget.

5.3 Final Evaluation with Simulator

At last, we evaluated our selected prescribed burn plan by re-running the simulations. First, we mark the fuel in the selected prescribed burn cells as 0, so they can no longer burn or spread fire further, and re-run the simulation to recompute the average wildfire impact metrics. With the same ignition points, we compare the average wildfire impacts without and with prescribed burns in Figure 9, and observe that the wildfire impact reductions are around 150 cells, 2000 forests, and 4000 for the degree of curing. Although these wildfire reduction values are different from the values from the surrogate function shown in Figure 5a, our solution does show success in reducing the wildfire impacts in all three objectives.

Table 2: The HYPERVOLUME results of different algorithms with different formulations for Dogrib, with 500 candidates. The hypervolume reported corresponds to the objective each formulation is optimizing.

	SMO-PB ($\times10^{10}$)	SMO-PB-F ($\times0.1$)	SMO-PB-SF ($\times10^{7}$)	SMO-PB-BF ($\times10^{10}$)
NSGA-II	5.36 ± 0.01	**9.56 ± 0.11**	1.67 ± 0.00	5.33 ± 0.00
NSGA-III	5.32 ± 0.02	9.46 ± 0.01	**1.68 ± 0.03**	**5.34 ± 0.01**
U-NSGA-III	5.32 ± 0.03	9.44 ± 0.01	1.67 ± 0.00	5.32 ± 0.00
C-TAEA	5.34 ± 0.03	8.94 ± 0.18	1.47 ± 0.05	5.31 ± 0.01
PPF	**5.38 ± 0.00**	-	-	-

Table 3: The results of NSGA-III with different formulations, reporting the number of Pareto front solutions (# Plans) for each formulation, and hypervolume with respect to the 3 main objectives only, the 3 fairness objectives only, and all 6 objectives (Dogrib with 500 candidates). SMO-PB-F and SMO-PB-SF result in as good hypervolume in the 3 main objectives as SMO-PB, while significantly improving the hypervolume in the 3 fairness objectives.

	# Plans	HPV(3 Main) ($\times10^{10}$)	HPV(3 Fairness) ($\times10^{-11}$)	HPV(3 Main + 3 Fairness) ($\times0.1$)
SMO-PB	21 ± 4	**5.32 ± 0.02**	1.24 ± 0.00	6.52 ± 0.02
SMO-PB-F	124 ± 31	5.31 ± 0.00	1.86 ± 0.00	9.46 ± 0.01
SMO-PB-SF	74 ± 6	5.30 ± 0.00	1.87 ± 0.04	**9.53 ± 0.18**
SMO-PB-BF	19 ± 4	2.82 ± 0.00	**1.90 ± 0.31**	4.85 ± 0.48

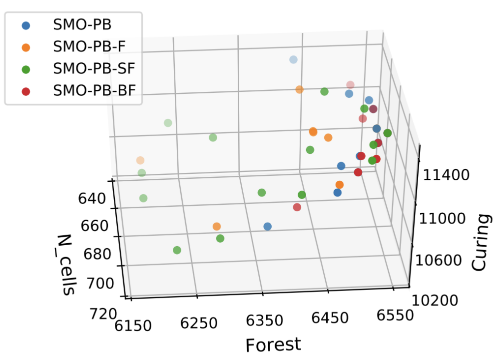

(a) The Pareto frontier of the formulations projected to the three main objectives.

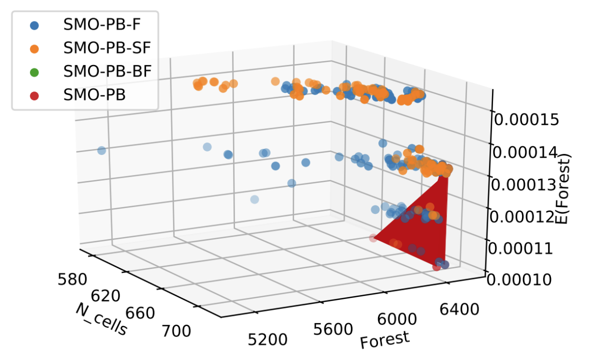

(b) The Pareto front of the formulations projected to the 'Area of cells', 'Forest burned' and 'Forest Fairness' objectives. We have specifically connected all points generated by SMO-PB as a surface.

Figure 5: The Pareto frontier of the formulations. The larger the value, the better the solution. Results for a single random seed.

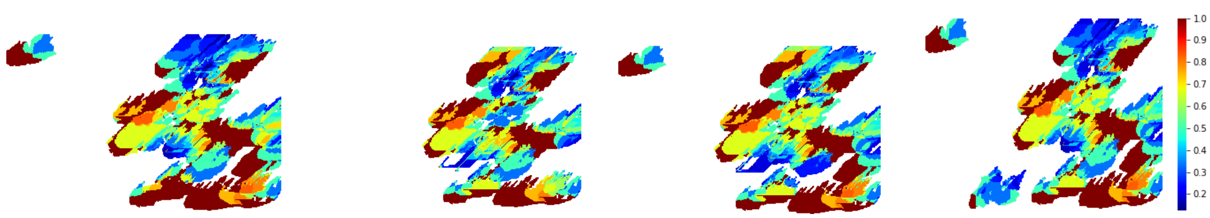

(a) SMO-PB **(b) SMO-PB-F** **(c) SMO-PB-SF** **(d) SMO-PB-BF**

Figure 6: The benefit heatmap of different formulation, corresponding to the points shown in Fig. 7.

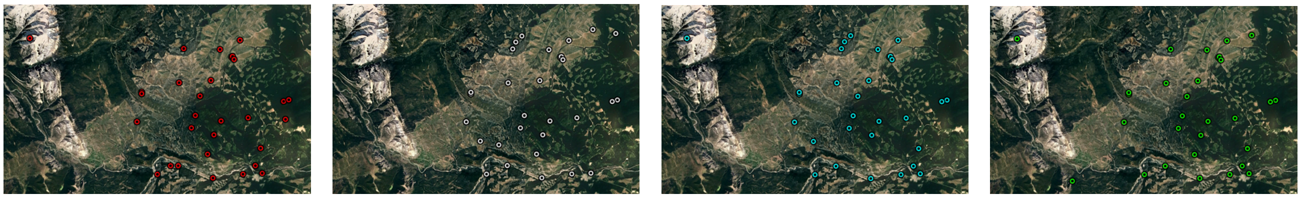

(a) SMO-PB (b) SMO-PB-F (c) SMO-PB-SF (d) SMO-PB-BF

Figure 7: One typical landscape prescribed burn plan from different formulations. The specific plan is randomly chosen from the Pareto frontier. Every algorithm has the same budget of 30. Most chosen units are shared in different formulations, but those candidates closer to the corner are only covered by the three variants with fairness concerns.

Table 4: The 3 main objectives hypervolume (HPV) results for NSGA-III with different number of candidates and different formulations. The bigger the number, the better the solution.

# Candidates	SMO-PB ($\times 10^{10}$)	SMO-PB-F ($\times 10^{10}$)	SMO-PB-SF ($\times 10^{10}$)	SMO-PB-BF ($\times 10^{10}$)
200	1.43 ± 0.01	1.42 ± 0.01	1.42 ± 0.02	0.79 ± 0.03
500	5.32 ± 0.02	5.31 ± 0.00	5.30 ± 0.00	2.82 ± 0.00
800	8.75 ± 0.02	8.48 ± 0.23	8.66 ± 0.05	4.43 ± 0.17

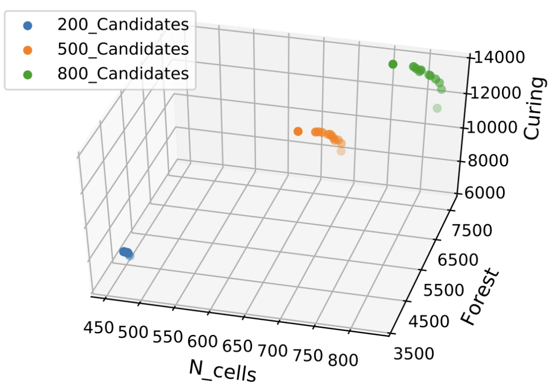

Figure 8: The Pareto frontier of ablation study with NSGA-III on different numbers of candidates. The larger the value, the better the solution.

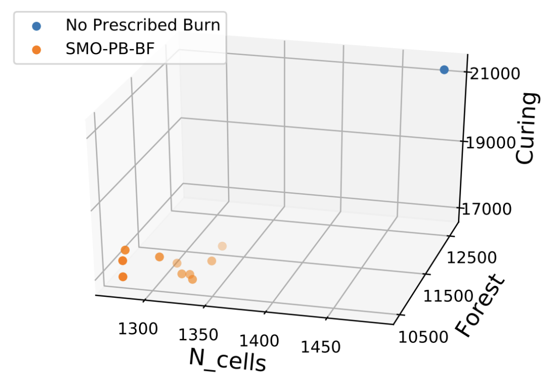

Figure 9: The average damage received the prescribed burn plan gives, calculated by re-running the simulator and calculating the weighted sum of burned cells. The smaller the number, the better the result.

6 CONCLUSION

In this work, we apply multi-objective optimization to solve the landscape-level prescribed burn planning for wildfire mitigation. We first give the general SMO-PB formulation, and then provide three formulation variants that introduce fairness into the problem to make the solutions provide a more spatially balanced distribution of prescribed burn benefits. We use a case study of Dogrib to validate that our formulations can yield a satisfying solution for a real world deployment, and show how introducing fairness can benefit more people without harming the overall benefit.

ACKNOWLEDGMENTS

This work was supported by NSF Award # KR 704735: "NSF Convergence Accelerator – Track D: Artificial Intelligence and Community Driven Wildland Fire Innovation via a WIFIRE Commons Infrastructure for Data and Model Sharing".

REFERENCES

[1] Soheila Aalami and Lina Kattan. 2017. Fair dynamic resource allocation in transit-based evacuation planning. *Transportation research procedia* 23 (2017), 400–419.
[2] Fermín J Alcasena, Alan A Ager, Michele Salis, Michelle A Day, and Cristina Vega-Garcia. 2018. Optimizing prescribed fire allocation for managing fire risk in central Catalonia. *Science of the total environment* 621 (2018), 872–885.
[3] Nikolaos Argyris, Özlem Karsu, and Mirel Yavuz. 2022. Fair resource allocation: Using welfare-based dominance constraints. *European Journal of Operational Research* 297, 2 (2022), 560–578.
[4] M Eugenia Captivo, Joao Climaco, Jose Figueira, Ernesto Martins, and Jose Luis Santos. 2003. Solving bicriteria 0–1 knapsack problems using a labeling algorithm. *Computers & Operations Research* 30, 12 (2003), 1865–1886.
[5] EA Catchpole, TJ Hatton, and WR Catchpole. 1989. Fire spread through non-homogeneous fuel modelled as a Markov process. *Ecological Modelling* 48, 1-2 (1989), 101–112.
[6] Carly Ruth Cowell and Chad Cheney. 2017. A ranking system for prescribed burn prioritization in Table Mountain National Park, South Africa. *Journal of*

environmental management 190 (2017), 283–289.
[7] Indraneel Das and John E Dennis. 1998. Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM journal on optimization* 8, 3 (1998), 631–657.
[8] Kalyanmoy Deb and Himanshu Jain. 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE transactions on evolutionary computation* 18, 4 (2013), 577–601.
[9] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
[10] Charles Delort and Olivier Spanjaard. 2010. Using bound sets in multiobjective optimization: Application to the biobjective binary knapsack problem. In *International Symposium on Experimental Algorithms*. Springer, 253–265.
[11] David Duran-Rodas, Dominic Villeneuve, Francisco C Pereira, and Gebhard Wulfhorst. 2020. How fair is the allocation of bike-sharing infrastructure? Framework for a qualitative and quantitative spatial fairness assessment. *Transportation Research Part A: Policy and Practice* 140 (2020), 299–319.
[12] Mark A Finney. 1998. *FARSITE, Fire Area Simulator–model development and evaluation*. Number 4. US Department of Agriculture, Forest Service, Rocky Mountain Research Station.
[13] Francis M Fujioka. 1985. Estimating wildland fire rate of spread in a spatially nonuniform environment. *Forest Science* 31, 1 (1985), 21–29.
[14] Pedro MM Guerreiro, Pedro JS Cardoso, and Hortênsio CL Fernandes. 2020. A comparison of multiple objective algorithms in the context of a dial a ride problem. In *International Conference on Computational Science*. Springer, 382–396.
[15] C Wayne Hanselka et al. 2009. Planning a Prescribed Burn. *Texas FARMER Collection* (2009).
[16] Rajendra K Jain, Dah-Ming W Chiu, William R Hawe, et al. 1984. A quantitative measure of fairness and discrimination. *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA* 21 (1984).
[17] Young-Hwan Kim, Pete Bettinger, and Mark Finney. 2009. Spatial optimization of the pattern of fuel management activities and subsequent effects on simulated wildfires. *European Journal of Operational Research* 197, 1 (2009), 253–265.
[18] Kathrin Klamroth and Margaret M Wiecek. 2000. Dynamic programming approaches to the multiple criteria knapsack problem. *Naval Research Logistics (NRL)* 47, 1 (2000), 57–76.
[19] Ke Li, Renzhi Chen, Guangtao Fu, and Xin Yao. 2018. Two-archive evolutionary algorithm for constrained multiobjective optimization. *IEEE Transactions on Evolutionary Computation* 23, 2 (2018), 303–315.
[20] Rodman Ray Linn, Scott L Goodrick, Sara Brambilla, Michael J Brown, Richard S Middleton, Joseph J O'Brien, and John Kevin Hiers. 2020. QUIC-fire: A fast-running simulation tool for prescribed fire planning. *Environmental Modelling & Software* 125 (2020), 104616.
[21] Thibaut Lust and Jacques Teghem. 2012. The multiobjective multidimensional knapsack problem: a survey and a new approach. *International Transactions in Operational Research* 19, 4 (2012), 495–520.
[22] Dmytro Matsypura, Oleg A Prokopyev, and Aizat Zahar. 2018. Wildfire fuel management: network-based models and optimization of prescribed burning. *European Journal of Operational Research* 264, 2 (2018), 774–796.
[23] The Department of Forestry and Fire Protection. [n. d.]. Wildfire Yearly Stats. http://www.fire.ca.gov/
[24] Cristobal Pais, Jaime Carrasco, David L Martell, Andres Weintraub, and David L Woodruff. 2019. Cell2fire: A cell based forest fire growth model. *arXiv preprint arXiv:1905.09317* (2019).
[25] Cristobal Pais, Jaime Carrasco, Pelagie Elimbi Moudio, and Zuo-Jun Max Shen. 2021. Downstream Protection Value: Detecting critical zones for effective fuel-treatment under wildfire risk. *Computers & Operations Research* 131 (2021), 105252.
[26] Annibale Panichella. 2019. An adaptive evolutionary algorithm based on non-Euclidean geometry for many-objective optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 595–603.
[27] Steven W Running. 2006. Is global warming causing more, larger wildfires? *Science* 313, 5789 (2006), 927–928.
[28] L Russo, P Russo, I Evaggelidis, and C Siettos. 2015. Complex network statistics to the design of fire breaks for the control of fire spreading. *Chemical Engineering Transactions* (2015).
[29] Haitham Seada and Kalyanmoy Deb. 2015. A unified evolutionary optimization procedure for single, multiple, and many objectives. *IEEE Transactions on Evolutionary Computation* 20, 3 (2015), 358–369.
[30] Ekunda Lukata Ulungu and Jacques Teghem. 1995. The two phases method: An efficient procedure to solve bi-objective combinatorial optimization problems. *Foundations of computing and decision sciences* 20, 2 (1995), 149–165.
[31] Yash Vesikar, Kalyanmoy Deb, and Julian Blank. 2018. Reference point based NSGA-III for preferred solutions. In *2018 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 1587–1594.
[32] Anthony LeRoy Westerling. 2016. Increasing western US forest wildfire activity: sensitivity to changes in the timing of spring. *Philosophical Transactions of the Royal Society B: Biological Sciences* 371, 1696 (2016), 20150178.
[33] J Wong, Jacques Sauve, and James Field. 1982. A study of fairness in packet-switching networks. *IEEE Transactions on Communications* 30, 2 (1982), 346–353.
[34] Qingfu Zhang and Hui Li. 2007. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on evolutionary computation* 11, 6 (2007), 712–731.
[35] Yiyang Zhao, Linnan Wang, Kevin Yang, Tianjun Zhang, Tian Guo, and Yuandong Tian. 2021. Multi-objective Optimization by Learning Space Partitions. *CoRR* abs/2110.03173 (2021). arXiv:2110.03173 https://arxiv.org/abs/2110.03173
[36] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. 2001. SPEA2: Improving the strength Pareto evolutionary algorithm. *TIK-report* 103 (2001).