

Applying Optimization Methods to Wildfire: A Simulation Study

Jacob W. Evans*, Burak Bagdatli†, and Dimitri N. Mavris‡
Aerospace Systems Design Laboratory, School of Aerospace Engineering
Georgia Institute of Technology, Atlanta, Georgia, 30332

Wildfire response is a wide field that spans many disciplines, from first response and fire thermodynamics to aircraft operations and system-of-systems analysis. While many fires can be addressed with only ground crews, fires in high-risk conditions (high temperature, low humidity, high wind, etc.) can be difficult to tame even for high-cost air-tankers. In such situations, aerial retardant drops are a critical tool, yet their effectiveness depends strongly on where and when they are deployed. Quantifying the quality of a retardant drop is therefore essential for guiding operational decisions and improving suppression outcomes.

Wildfire growth, however, is inherently a multi-objective problem. It involves tradeoffs between minimizing burned area, reducing risk to human populations, limiting emissions, and slowing fire spread. This complexity makes it difficult to define a single measure of success and to apply optimization techniques directly. To address this challenge, this work compares multiple novel cost functions to compress one or more spatial variables related to wildfire growth into a single scalar variable representing the “level of undesirability” of future fire growth. The cost functions integrate metrics of area burned, fire progression, population exposure, or emissions, thereby capturing both local and global impacts of suppression strategies.

Using these cost functions, an optimization of retardant drop locations for aircraft operations was conducted using brute force search, Bayesian optimization, and differential evolution. Results demonstrate that the proposed framework can systematically evaluate tradeoffs between competing objectives and provide actionable guidance for aerial firefighting operations. Beyond immediate tactical applications, this approach offers a foundation for integrating ecological, social, and environmental priorities into wildfire management, supporting more resilient and sustainable suppression strategies.

I. Nomenclature

a	=	cell area
b	=	burned cells
c	=	minimum travel time matrix
m	=	simulation time
s	=	scalar cost function (fire score)
p_{norm}	=	normalized population density
c_{norm}	=	normalized minimum travel time matrix
\odot	=	hadamard product
e_{CO_2}	=	CO ₂ emissions from burned area

II. Motivation

WILDLAND fire is a very complicated topic requiring significant research [1]. Scott [2] gives a very thorough description of the relevant definitions of wildfires. Fire is a “self-perpetuating combustion characterized by the emission of heat and accompanied by flame and/or smoke” [2]. It requires heat, fuel, and oxygen to maintain combustion. A wildfire, often more accurately referred to as a vegetation fire or landscape fire, is a specific type of fire that differs from structure fires or other human-made fuels. Different types of vegetation burn with different characteristics. Previous

*Ph.D. Candidate, Graduate Researcher, jevans97@gatech.edu

†Research Engineer II, burak.bagdatli@ae.gatech.edu, AIAA Member

‡Georgia Tech Distinguished Regents Professor and Director of ASDL, AIAA Fellow

research has attempted to discretize the variability of vegetation in a set of fuel models [3–5]. If a fire is burning in a flat area where vegetation is approximately spatially constant, the fire front can be approximated as an ellipse [6, 7] in the direction of the local wind vector. However, in a realistic landscape, spatially varying fuel models combine with local wind and slope to create a very complex external fire front. Additionally, certain areas of the landscape are more important than others to protect, either because they hold environmental value (e.g., a historical forest), monetary value, (e.g., a home or town), or strategic value (e.g., a choke point). The combination of value and complexity creates an environment ripe for computational analysis and optimization.

It should be of little surprise that certain areas of a landscape are more important to protect than others. Examples of critical areas include residential homes on a Wildland Urban Interface (WUI) [8], valleys that are key to growth, roadways that are key evacuation routes, etc. Ideally, fire response would protect all critical areas and extinguish the fire very quickly, but as the wildfire increases in size and complexity, trade-offs must be made. The highest priority is firefighter and public safety [9], followed by structures. Structures in the WUI are of particular importance to wildland fire response in the American West, because “half of the residents [. . .] live in the wildland-urban interface” [10]. The cost of relocation is approximately \$2500 per person [11, 12].

After protecting lives and homes, the next target is the regions of the landscape that provide the highest return on investment. A C-130 can cost over \$6000 per hour to operate and a 5000 gal retardant drop can cost an additional \$15,000 [13]; therefore, maximizing the efficiency of each retardant drop is essential. In addition, a well-placed drop can significantly reduce the amount of land burned over the course of a whole fire, which can significantly reduce restoration costs. Historically, the process of deciding where a retardant drop should fall spatially was determined by the incident commander using their previous knowledge, but since fire conditions vary from fire to fire, this process does not guarantee an optimal solution.*

One of the ways that a fire can vary is the wind condition. Wind is one of the highest contributors to fire growth speed [3, 14], and the wind can change at any time. The wind also can vary spatially as shown in Fig. 1. While a wide range of methods can be used to describe the current state of a vegetation fire, predicting future growth is challenging. The method used for modeling future wildfire growth in this work has the ability to vary the wind conditions in space and time, allowing for the complicated wind effects to be included.

III. Background

Contrary to common media perceptions, the vast majority of wildfires are relatively tame. Of the 3 to 4.6 million km² that burn every year (approximately 4% of the global land surface) [16], just 3% of these fires drive over 80% of associated fire damages [17]. For such extreme fire events, the best that can be expected of firefighters is to slow down the fire, remove people from the fire’s path, and clean up afterwards. The method for slowing down a fire relies largely on ground crews physically removing fuel in lines to slow or stop growth in key areas combined with retardant and water drops that increase the moisture and/or reduce the flammability of the fuel. The current method for choosing the location to place water or retardant or a fire line is largely based on the incident commander’s intuition combined with some rules of thumb [18]. Those rules of thumb have been supported in part the computational wildfire modeling, which has progressed from 1-D mathematical relationships [14] to full growth models [3, 19]. This work seeks to add a level of optimization-based intelligence on top of the wildfire growth by asking the question: “If a retardant drop were placed here, what would the resulting growth look like?” Then, by varying the retardant drop location, the growth will change. By placing the location of the retardant drop in the control of an optimizer, the optimal drop location can be found. The goal of this work is to enhance firefighters’ ability to combat high-risk fires by using computational analysis to optimize the placement of retardant drops. This is accomplished in a wildfire simulation environment.

A. Wildfire Modeling

The wildfire modeling environment used in this work is the Minimum Travel Time (MTT) algorithm used in FlamMap from Missoula Fire Sciences Lab [19–22]. FlamMap takes an $m \times n \times 8$ matrix of topological information of a specified region. The horizontal grid size is given by the integers m and n . The eight layers give spatial elevation, slope, aspect, fuel model, canopy cover, canopy height, crown base height, and crown bulk density, respectively. It also requires an ignition point or polygon and an input file that describes other inputs such as the wind conditions and simulation time. From these inputs, FlamMap returns a series of files describing the future growth of the fire during the prescribed simulation time. These files include the flame length at each cell, the location of ember drops (for spotting),

*In-person interview with Thomas Shoots, CalFire, June 2023

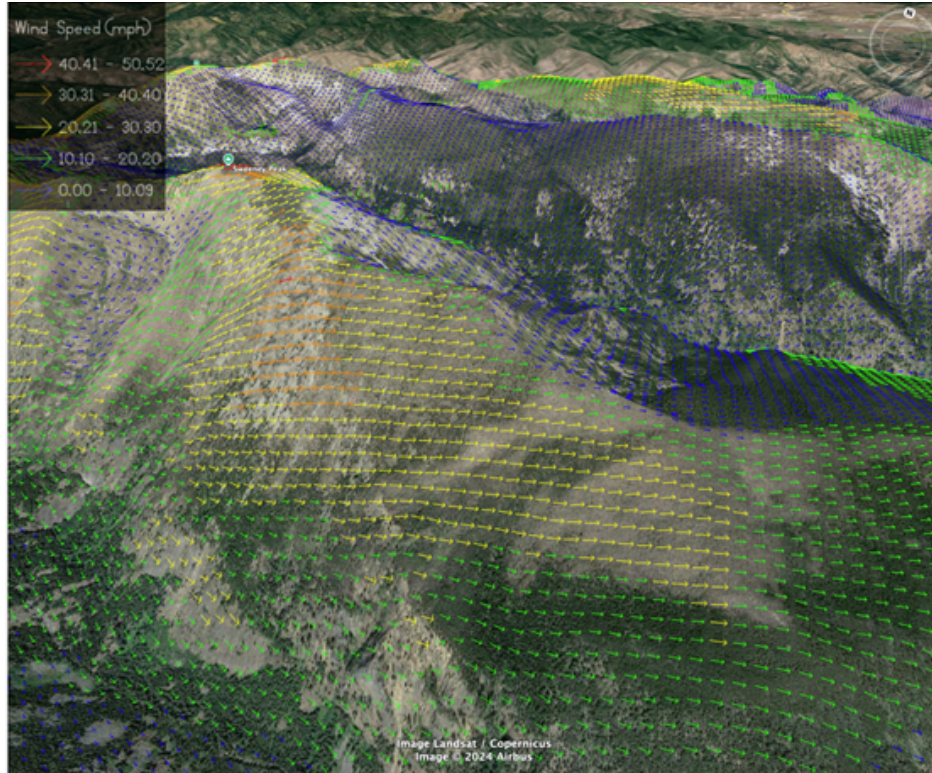


Fig. 1 Spatial variation of wind in Happy Camp, CA as calculated by wind ninja [15] for an input average wind of 10 mph at 270 degrees

the rate of spread for each cell, the influence of each cell, and so on. The only output used for this analysis is the Arrival Time (the time after ignition that the fire front arrives at that cell). The arrival time matrix (also called the minimum travel time matrix) is the same dimensions as one layer of the input topological matrix ($m \times n \times 1$).

B. Optimization

The goal of optimization is to minimize an objective function, typically within some constraints, by changing a range of possible input parameters [23]. Depending on the form of the objective function, a range of optimizers can be used. Broadly, optimizers can be classified into several categories:

Gradient-based methods: These rely on derivatives of the objective function to guide the search (e.g., steepest descent, Newton’s method). They are efficient for smooth, differentiable problems but cannot be applied when gradients are unavailable or unreliable.

Gradient-free deterministic methods: These include direct search techniques such as grid search or pattern search. They are simple to implement but often scale poorly with dimensionality.

Stochastic methods: Algorithms such as genetic algorithms, simulated annealing, and differential evolution introduce randomness to explore the search space. They are robust to local minima but can require significant computational effort.

Probabilistic model-based methods: Bayesian optimization and related surrogate-model approaches build statistical models of the objective function and use them to guide sampling. These are particularly effective when evaluations are expensive but limited in number.

For spatial wildfire growth as a function of retardant drop location, there is no gradient (first derivative) or Hessian (second derivative) available, so gradient-based methods must be excluded. For this study, three optimizers were selected, one from each of the remaining three categories, to compare their respective ability to find an optimum retardant drop location:

Brute force search: A gradient-free deterministic method, brute force sets a baseline by exhaustively evaluating candidate solutions. While computationally expensive, it ensures that the global optimum is found within the

discretized search space.

Bayesian optimization: A stochastic method, Bayesian optimization offers sample efficiency by constructing a surrogate model of the objective function and selecting new points based on expected improvement. This is well-suited to problems where each evaluation (e.g., simulating wildfire spread) is costly.

Differential evolution: A population-based stochastic algorithm that perturbs candidate solutions through recombination and mutation. It is effective in exploring complex, multimodal landscapes and avoids premature convergence.

C. Background on Cost Functions

Optimization in wildfire management requires the definition of cost functions that translate physical fire behavior into quantifiable objectives. These functions serve as the link between simulation outputs and decision-making criteria, allowing different strategies to be compared on a consistent basis. In this study, four cost functions were developed to capture distinct dimensions of wildfire impact: burned area, fire intensity, population exposure, and emissions.

The first cost function, area burned, is a direct measure of fire extent. It reflects the most fundamental objective of suppression efforts—reducing the total land affected by fire. While simple, this metric does not account for variations in land use or population density, and therefore provides only a coarse measure of impact.

The second cost function, the simple fire score, incorporates fire arrival time to provide a measure of fire severity across the landscape. This metric emphasizes the temporal dynamics of fire spread, rewarding strategies that delay or slow fire progression. However, it remains spatially uniform and does not differentiate between areas of higher or lower value.

To address this limitation, a third cost function was introduced: the population-weighted fire score. This metric prioritizes protection of human communities. It reflects the social dimension of wildfire management, where minimizing exposure and risk to populated areas is often the highest priority. The tradeoff is that ecological or remote areas may be undervalued in this formulation.

Finally, the emissions-based fire score quantifies the environmental cost of wildfire in terms of carbon dioxide released. This metric highlights the role of wildfires in contributing to greenhouse gas emissions and climate change. Unlike the other metrics, it directly connects local fire suppression decisions to global environmental impacts. The challenge is that emissions estimates depend on fuel type and combustion completeness, introducing uncertainty into the calculation.

IV. Methodology

This work uses Happy Camp, California area, shown in Fig. 2 as a case study due to its proximity to the Happy Camp Complex Megafire that burned over 83,000 acres and cost over \$47.4 million [24, 25]. The Happy Camp region is an example of WUI and provides a significant variation in fuel types as shown in Fig. 3, with 27 fuel types used out of a possible 40, making it an ideal case study.

A fire is instantiated at the midpoint of the map and allowed to grow for 5000 minutes. At the end of the simulation time, the Arrival Time matrix is returned. A small scale example of how an MTT might look is show in Table 1. The *not a number (nan)* entries around the perimeter indicate cells that do not catch fire, either because they contain an unburnable fuel model or because the fire would not reach them until after the simulation time had expired. The colors in this example represent the growth of the fire, where red shows where the fire currently is (cells with 0) and green shows cells that will take the longest time to catch fire. The green row on the far left and the orange on the middle-right implies that the fire will grow fastest to the right.

While this example is only 120 cells, the output from FlamMap for the Happy Camp case study contains 741,600 cells of potential arrival times, far too many to optimize individually. A cost function combines all these cells into a single scalar value that represents not only the area that the fire burns, but the speed at which it grows. The equation for this cost function is given by Eq. 1, where s is the scalar cost function (fire score), c is the MTT matrix, m is the simulation time, and a is the cell area.

$$s = \frac{\sum(m - c)}{ma} \quad (1)$$

A retardant drop is simulated by modifying the fuel model of the cells that were dropped upon. The modification can be either a perfect drop (modifies the cells to an unburnable fuel type like water) or an imperfect drop (modifies the cells to a custom fuel that has a higher moisture content than the original, which slows down growth but does not stop it). The location of a retardant drop has, in general, three degrees of freedom: latitude of midpoint, longitude of midpoint,



Fig. 2 Happy Camp Complex Fire from [25]




Fig. 3 Fuel model in the Happy Camp region. Blue represents water (unburnable), grays represent urban/suburban development (unburnable), greens are generally slow-moving fuels, and yellows are generally fast-moving

Table 1 Minimum Travel Time example

nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
nan	65	47	42	43	35	32	40	34	nan
nan	62	22	25	17	26	25	20	28	nan
nan	60	17	0	0	11	9	15	18	nan
nan	55	0	0	0	0	8	12	15	nan
nan	53	0	0	0	0	0	6	11	nan
nan	54	0	0	0	0	0	5	12	nan
nan	59	45	0	0	0	8	13	20	nan
nan	84	50	0	0	15	18	20	22	nan
nan	102	73	26	28	24	26	28	35	nan
nan	123	69	48	53	48	43	39	46	nan
nan	nan	nan	nan	nan	nan	nan	nan	nan	nan

and angle of drop with respect to north. However, a drop angle perpendicular to the contour lines of the fire, as shown in Fig. 4a, will provide negligible benefit; therefore, the angle degree of freedom can be removed by assuming that the retardant drop would always be parallel to the fire contour at the midpoint of the drop location, as shown in Fig. 4b. Also, a retardant drop closer to the perimeter of the fire will be better than farther, so another degree of freedom was removed by prescribing that the midpoint of the retardant drop would be on the fire perimeter at the time that the aircraft arrived. This can be seen as the highlighted black contour line of Figures 4a and 4b. The final drop logic contains only a single degree of freedom: where on the perimeter contour of the fire to place the midpoint of the retardant drop.

Once a retardant drop is created on a map and the fuel model is changed due to the drop, FlamMap (or any other fire model) can be re-run with the updated fuel model to predict the future fire growth due to that retardant drop. Any variation in drop location will result in a slightly different MTT grid, which will result in a slightly different fire score. This change in fire score can be recognized by an optimizer to find the optimal retardant drop location.

A few limitations exist for the optimization process. First, because the fire score is a function of every cell in the map, there is no gradient or Hessian available. Second, it has been found in preliminary testing that the variation of the fire score for a given perimeter curve is very noisy, which eliminates simple 1-D minimizer like Golden Section or Brent’s algorithm [26] and reduces the likelihood of success for a local optimizers. For this reason, this study focuses on optimizers that are gradient-free and are intended to find a global minimizer among many possible local minima (of course the only way to guarantee a global minimizer is to evaluate every possible input, which is impossible on a continuous input space).

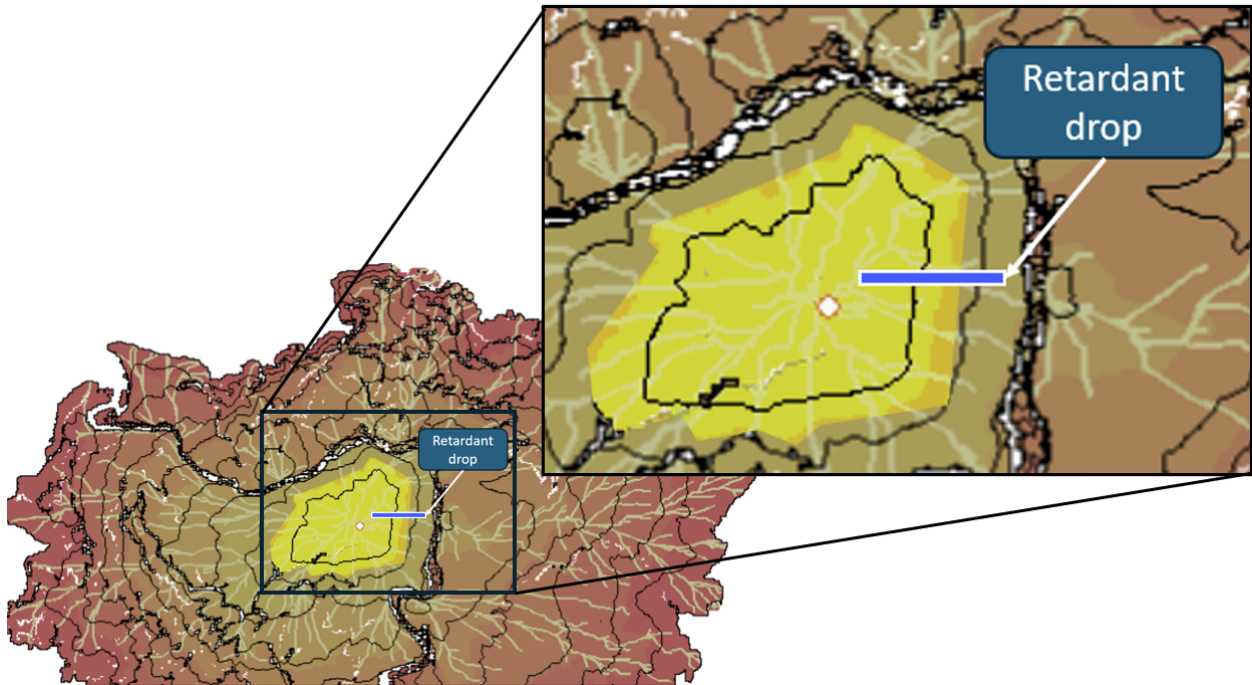
A. Cost Function Definitions

Four cost function metrics were defined for the purpose of this study. Area burned is simply the predicted area that will be burned in the simulation time. It is calculated as the number of cells in the arrival time matrix that are not NaN multiplied by the cell area as given in Eq. (2). In the case of FlamMap, the cell area is a constant of $900m^2$. The next metric is the simple fire score defined in Eq. (1). The third metric is a fire score that includes population as a factor. It uses the a matrix of normalized population density multiplied by the normalized fire arrival time matrix as shown in Eq. (3). The final metric tested was an emissions-based fire score. This score returns the total pounds of CO_2 emitted by the fire that burns within the simulation time as given in Eq. (4).

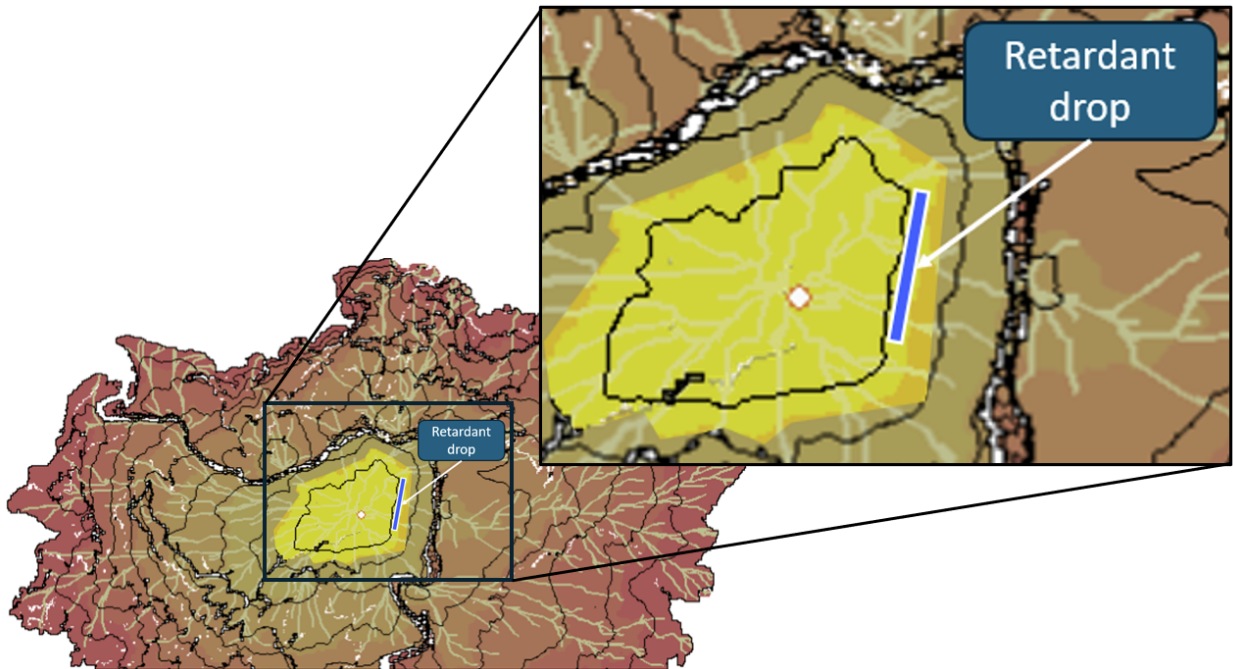
$$A = b \cdot a \tag{2}$$

$$S_{population} = \frac{\sum(c_{norm} \odot p_{norm})}{a} \tag{3}$$

$$S_{emissions} = \sum e_{CO_2} \tag{4}$$



(a) Retardant drop perpendicular to fire contour does not provide much value in slowing fire growth



(b) Parallel retardant drop to fire contour is higher quality

Fig. 4 Effects of dropping retardant at different angles to the fire front.

B. Optimization Functions

Three optimization functions were evaluated in this study. The first is the simple brute force optimization given in Algorithm 1. In brute force optimization, the cost function is evaluated at known intervals and the minimum value is returned as the optimum. This is similar in principle to a full factorial Design of Experiments. In the example at hand, a retardant drop was evaluated at 10 degree intervals around perimeter of the fire.

Algorithm 1: Brute Force Fire Score Optimizer

Input: Set of angles A
Output: Best angle θ^* and maximum fire score S^*

```
1  $S^* \leftarrow -\infty$ ;  
2  $\theta^* \leftarrow \text{null}$ ;  
3 foreach  $\theta \in A$  do  
4    $S \leftarrow \text{EvaluateFireScore}(\theta)$ ;  
5   if  $S > S^*$  then  
6      $S^* \leftarrow S$ ;  
7      $\theta^* \leftarrow \theta$ ;  
8   end  
9 end  
10 return  $\theta^*, S^*$ ;
```

The next optimizer that was evaluated was the Bayesian optimization as detailed in Algorithm 2. A Bayesian optimization is most useful where the cost function is a black box (i.e., its values and gradients are unknown *a priori*), is expensive to evaluate, or may be noisy.

In the case of a fire score as a cost function, all three criteria are met. The relationship between retardant drop location and resulting fire score is a black box due to the spatial nature of the resulting value. Computation of the fire score requires calculation of the minimum travel time matrix for each retardant drop location. And while this study is using settings that ensure the fire score is deterministic, spotting, changing winds, and changing moisture levels can all add stochasticity to the resulting metrics.

Algorithm 2: Bayesian Optimization

Input: Set of angles A , objective function $\text{EvaluateFireScore}(\theta)$
Output: Best angle θ^* and maximum fire score S^*

```
1 Initialize Gaussian Process surrogate model;  
2 Initialize acquisition function (e.g., Expected Improvement);  
3 while budget not exhausted do  
4    $\theta \leftarrow \text{SelectNextPoint}(A, \text{surrogate}, \text{acquisition})$ ;  
5    $S \leftarrow \text{EvaluateFireScore}(\theta)$ ;  
6   Update surrogate model with  $(\theta, S)$ ;  
7 end  
8  $\theta^*, S^* \leftarrow \text{BestObservedPoint}()$ ;  
9 return  $\theta^*, S^*$ ;
```

Finally, the third optimizer evaluated was the Differential Evolution (DE) algorithm shown in Algorithm 3. Differential evolution is a population-based, evolutionary algorithm particularly well-suited for continuous optimization problems. Unlike brute force, which exhaustively evaluates all candidates, or Bayesian optimization, which builds a surrogate model, DE maintains a population of candidate solutions and iteratively improves them through mutation, crossover, and selection. This allows DE to balance exploration of the search space with exploitation of promising regions. DE is robust to noisy, non-linear, and multi-modal cost functions, making it a strong candidate for optimizing fire score in complex fire spread scenarios.

The three optimization approaches each offer distinct advantages and trade-offs. Brute force optimization is simple, transparent, and guarantees identification of the global optimum within the evaluated grid, but it quickly becomes computationally expensive as the search space grows. Bayesian optimization, by contrast, is highly efficient for expensive

Algorithm 3: Differential Evolution Fire Score Optimizer

Input: Objective function $\text{EvaluateFireScore}(\theta)$, population size N , mutation factor F , crossover rate CR
Output: Best angle θ^* and maximum fire score S^*

- 1 Initialize population $P = \{\theta_1, \theta_2, \dots, \theta_N\}$ randomly from A ;
- 2 Evaluate fire scores $S_i = \text{EvaluateFireScore}(\theta_i)$ for all i ;
- 3 **while** *budget not exhausted* **do**
- 4 **foreach** $\theta_i \in P$ **do**
- 5 Select three distinct candidates $\theta_a, \theta_b, \theta_c$;
- 6 $\theta_{mut} \leftarrow \theta_a + F \cdot (\theta_b - \theta_c)$;
- 7 $\theta_{trial} \leftarrow \text{Crossover}(\theta_i, \theta_{mut}, CR)$;
- 8 $S_{trial} \leftarrow \text{EvaluateFireScore}(\theta_{trial})$;
- 9 **if** $S_{trial} > S_i$ **then**
- 10 $\theta_i \leftarrow \theta_{trial}$;
- 11 $S_i \leftarrow S_{trial}$;
- 12 **end**
- 13 **end**
- 14 **end**
- 15 $\theta^*, S^* \leftarrow \text{BestCandidate}(P)$;
- 16 **return** θ^*, S^* ;

black-box functions: it leverages a surrogate model to intelligently balance exploration and exploitation, often finding near-optimal solutions with far fewer evaluations. However, its performance depends on the quality of the surrogate model and acquisition function. Differential evolution provides a robust, population-based search that is well-suited for rugged, non-linear, and multi-modal landscapes. It does not require gradient information and can escape local optima, but may require careful tuning of parameters such as mutation factor and crossover rate. Together, these methods illustrate a spectrum of optimization strategies: exhaustive search (brute force), model-guided search (Bayesian), and evolutionary search (DE).

C. Case Study

This work features two primary fires, defined here as Fire 1 and Fire 2. They differ only in the ignition point location and the wind speed and direction. The map, which remains constant across all cases, consists of the 8 layers around the town of Happy Camp, CA. The fuel model ranges from fuel 91 (Urban/Developed) to fuel 189 (Very High Load Broadleaf Litter) [27]. The map consists of 800×927 cells that are 30 m-by-30 m for a total area of 164,928 acres. The fuel moisture is assumed to be 3% for 1 hour fuel, 4% for 10 hour fuel, 7% for 100 hour fuel, 49% for live herbaceous, and 82% for live woody. This scenario is between "low" and "very low" moisture from tables 3 and 4 for [27]. The population map is provided by [28]. This area of Wildland-Urban Interface is characterized by a very low population count, with a little over 1000 people living on the entire map and no more than four people per 30×30 meter cell. The emissions data is calculated using Spatial First Order Fire Effects Model (FOFEM) [29]. While many gas emissions results were available, CO_2 was chosen for this study.

Fire 1 has an ignition in the center of the map (0,0). It has a wind direction from due north with a speed of 10 miles per hour, which blows the fire toward the populated areas south of the river. Two scenarios of Fire 1 were analyzed. The first, to be called Fire 1A, had a simulation time of four days (5760 minutes), and the second, to be called Fire 1B, had a simulation time of eight days (11,520 minutes). Fire 1A is shown in Fig. 5, where each black contour line is the fire perimeter after one day (1440 minutes).

Fire 2 has an ignition at (6000, -2000) m in map coordinates, which is just north of the town of Happy Camp. It has a wind direction of 120 degrees clockwise from due north, blowing from the southeast to the northwest, away from the population, at 20 miles per hour.

For all cases, the retardant drop location is on the 2880 minute contour, simulating a scenario where a retardant drop is conducted exactly two days after fire ignition. The simulation time being longer than two days (four and eight days, respectively) allows the impact of that drop to affect the predicted future growth of the fire.

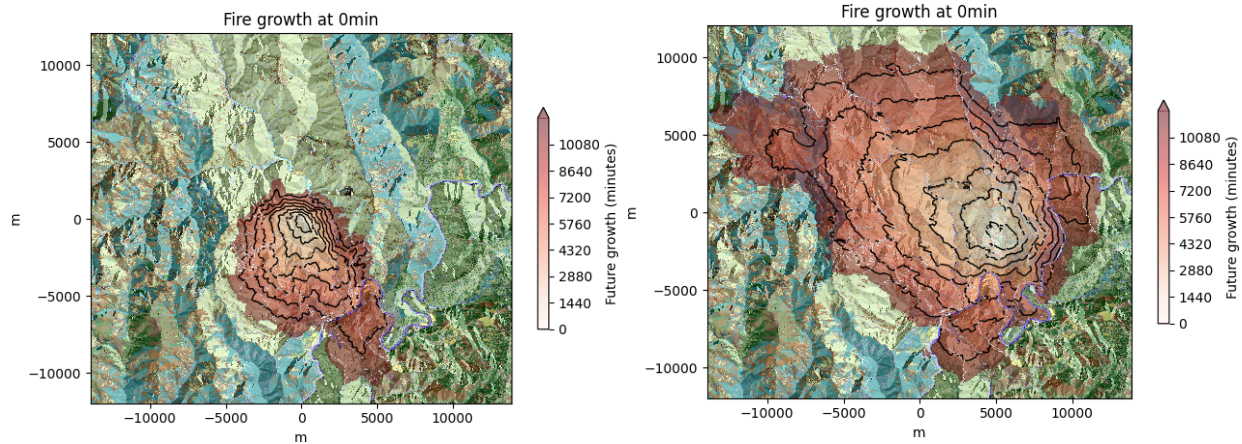


Fig. 5 Fire 1A/B (left) and Fire 2A/B (right) future growth with no retardant drops. The B version (eight-day simulated fire) is shown while the A version (four-day) is the first four contour lines (up to minute 5760)

V. Results and Discussion

The optimization experiments were conducted using three gradient-free methods: brute force search, Bayesian optimization, and differential evolution. In theory, all three methods should converge to the same optimal drop location (the global minimum) for each metric. The global minimum for each metric may or may not be at the same location. It was found that, instead of a single global minimum, a handful of local minima were found by the three optimizers for different cases. For Fires 1A and 1B, where the fire ignition point is far from populated areas and is being pushed by the wind toward population, one would expect the optimal drop location for all metrics to be to the southeast. Fig. 6 shows this to be mostly true. For Fire 1B, all but three of the optimal drop locations are to the southeast. For Fire 1A, there is a wider gap between the found local minima. This means that increasing the simulation time from four days to eight days does have an impact on the optimality of the solution.

For Fires 2A and 2B, shown in Fig. 7, the difference is even more pronounced. The four day simulation shows optimizer convergence on all sides of the fire, while the eight day simulation shows convergence at two primary locations. All but four cases for Fire 2B converged to the northwest, and the remaining cases converged due north of the ignition point. Once again this implies that a longer simulation time improves the performance of all four methods. It is interesting that the population score gave the same optimal drop location as area burned and fire score. There are a few possible explanations: first, the number of people on the whole map was too small to have a meaningful impact on directionality; second, the population density on the town of Happy Camp was so small compared to the population of the surrounding areas such that the population layer was effectively identical to a multiplying all cells of the burned area by a constant; third, there is so much noise that the optimizer finding the same location for all three population cases is coincidence only. In order to understand how each optimizer arrived at its optimal drop location, let us examine the convergence progression for each case over time.

A. Comparison of Optimizer Performance

Figure 8 shows the current best value for each iteration for each metric to until convergence for Fire 1A, while Figure 9 shows the same for Fire 1B. If a new function evaluation tested by the optimizer is a lower value than the previous best, the plot will decrease to the new best. Otherwise it will remain constant. One way to read it is “if the optimizer were hit its maximum iteration limit at this iteration, what would the function value be”. For Fire 1A and 1B, the Bayesian optimization consistently identified lower optimal values of the cost function compared to brute force and differential evolution and in far fewer function evaluations. This outcome suggests that the surrogate-model approach is more effective at exploiting promising regions of the search space, even when the objective function is complex and non-differentiable. Brute force guarantees exhaustive coverage but is limited by discretization. This allowed it to get close to the optimum found by the other methods but never the best. Differential evolution provides robust exploration but is limited by reaching its stopping criterion early. DE found the best point for Fire 1A—Area Burned, but for the other seven cases it performed poorly.

In contrast, for Fire 2A, shown in Fig. 10, brute force consistently found the best point over the other two optimizers.

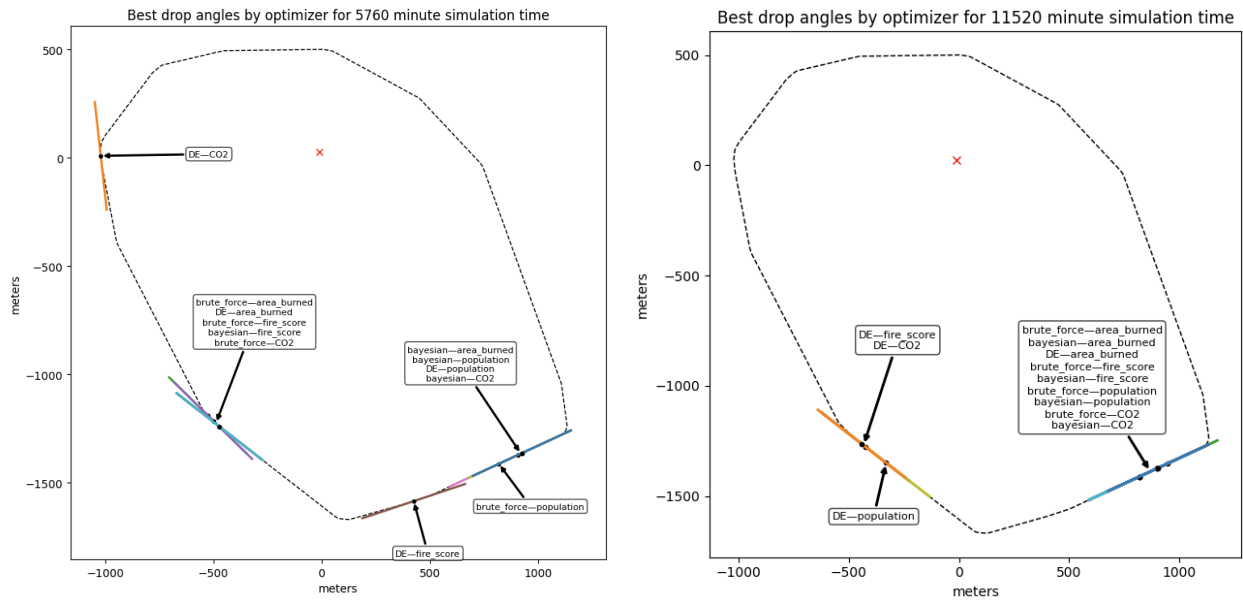


Fig. 6 Best retardant drop location defined by each optimization method and cost function metric for Fire 1A (four-day simulation time) on the left and 1B (eight-day simulation time) on the right. Dashed line is the design space

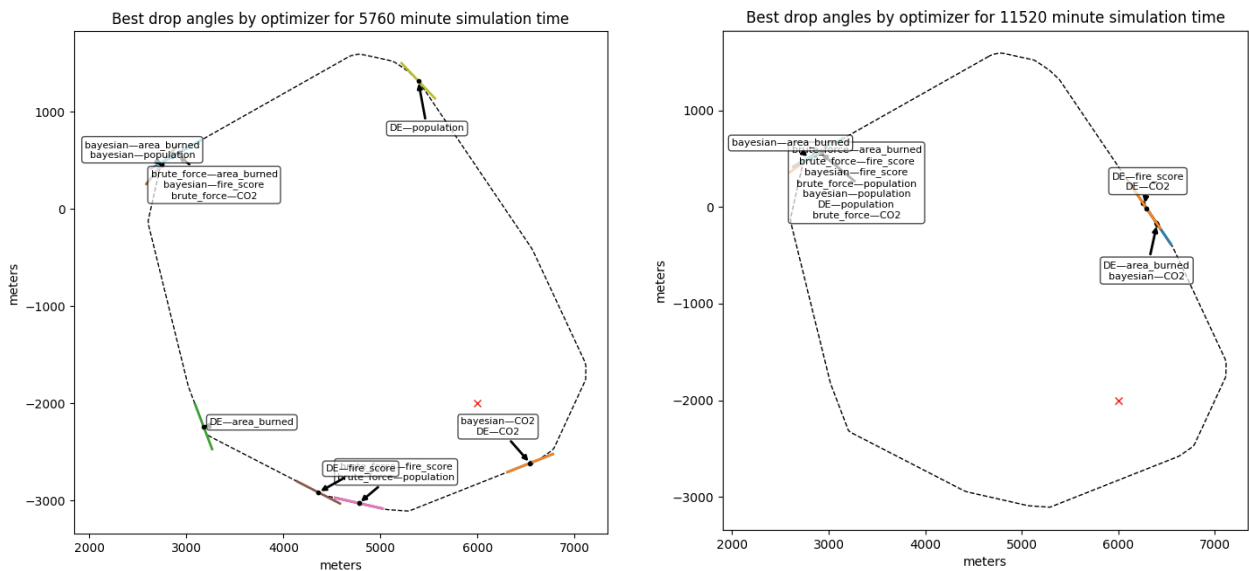


Fig. 7 Best retardant drop location defined by each optimization method and cost function metric for Fire 2A (four-day simulation time) on the left and 2B (eight-day simulation time) on the right. Dashed line is the design space

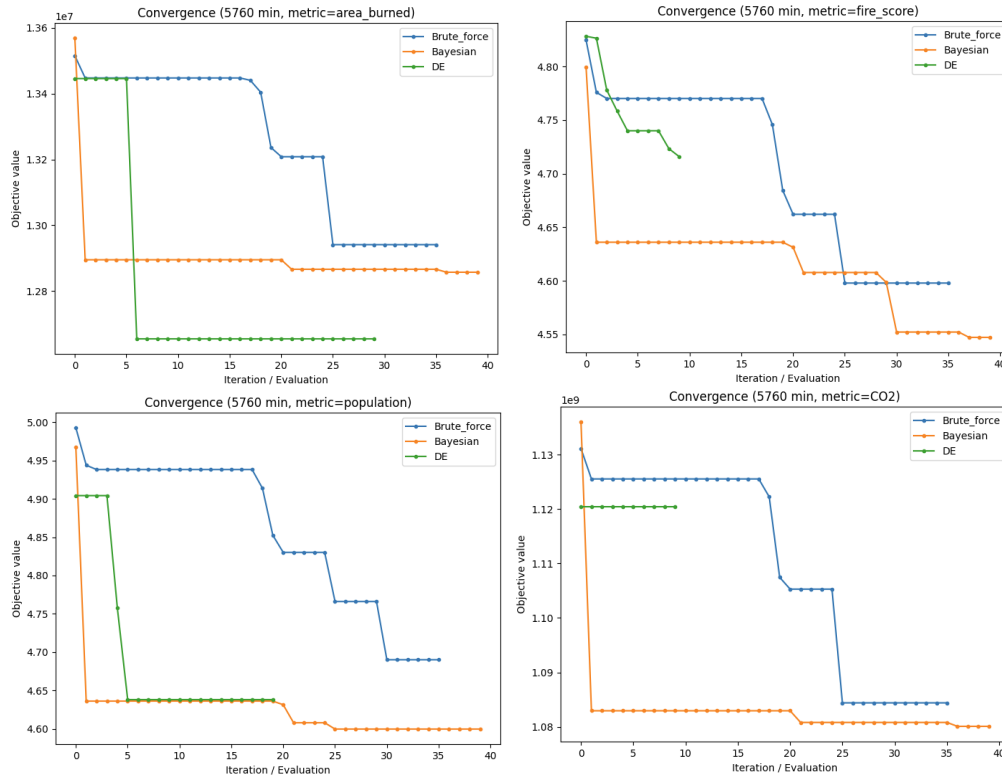


Fig. 8 Optimum value per iteration for all four metrics for Fire 1A

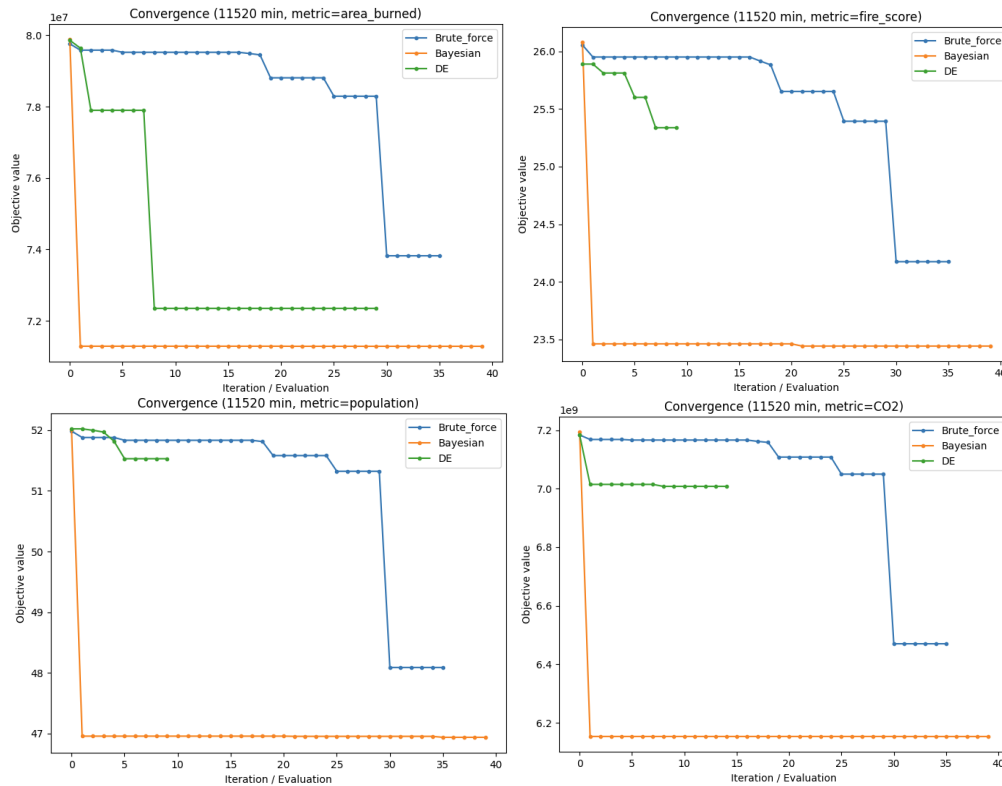


Fig. 9 Optimum value per iteration for all four metrics for Fire 1B

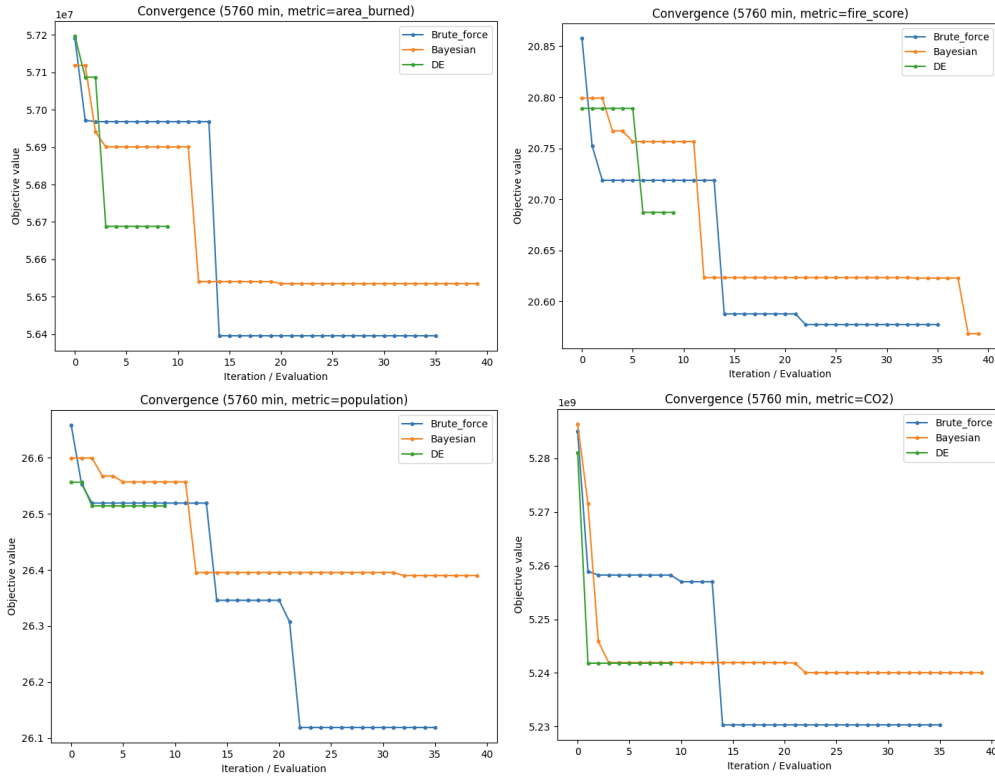


Fig. 10 Optimum value per iteration for all four metrics for Fire 2A

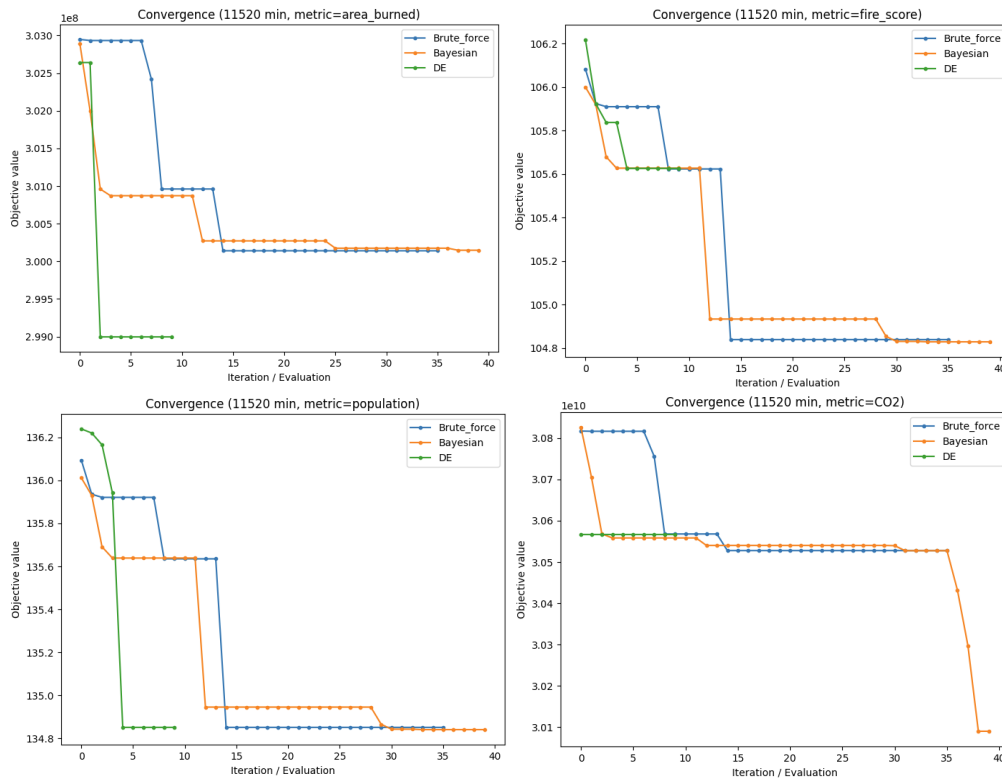


Fig. 11 Optimum value per iteration for all four metrics for Fire 2B

Table 2 Run times (minutes:seconds) for Fires 1A and 1B across optimizers and metrics

Sim Time (min)	Metric	Brute Force	Bayesian	DE
5760	Area Burned	7:22	8:14	7:19
5760	Fire Score	7:22	8:14	2:15
5760	Population	7:22	8:15	4:42
5760	CO ₂	7:22	8:14	2:15
11520	Area Burned	50:11	55:20	48:35
11520	Fire Score	50:20	54:24	15:16
11520	Population	51:12	53:56	15:24
11520	CO ₂	50:15	53:56	23:19

Table 3 Run times (minutes:seconds) for Fires 2A and 2B across optimizers and metrics

Sim Time (min)	Metric	Brute Force	Bayesian	DE
5760	Area Burned	31:33	35:09	9:36
5760	Fire Score	31:33	34:58	9:40
5760	Population	31:29	34:57	9:44
5760	CO ₂	31:29	34:58	9:36
11520	Area Burned	244:17	269:56	74:20
11520	Fire Score	244:05	270:20	74:52
11520	Population	243:53	270:11	74:28
11520	CO ₂	244:12	270:44	74:44

In all four of these cases, differential evolution reached its stopping criterion early and provided the works objective value. Fire 2B is significantly more difficult to identify a “winner.” For the area burned metric, differential evolution found the best value in the fewest iterations. For fire score, brute force and Bayesian tied for first. For population score, it was a three-way tie. And for emissions, Bayesian found the best value but took the most iterations to do so. Interestingly, looking back at the right side of Fig. 7, the area burned optimum found by differential evolution and the emissions optimum found by Bayesian are the outliers due north of the of the fire, not the cluster to the northwest found by the majority of the optimizers. This means that while the northwest is intuitive drop location because it is downwind and was identified in the majority of cases, it is actually only a local minimum for at least the area burned and emissions metrics, while the global minimum (of the points tested) is the counterintuitive point to the north that appeared to be an outlier.

B. Runtime Tradeoffs

Table 2 reports the average runtime for each optimizer across the tested metrics for Fires 1A and 1B, while Table 3 reports the same for Fires 2A and 2B. In all 16 cases, the brute force and Bayesian optimizers showed comparable run times. This is because neither of these methods included a stopping criterion, so they ran a constant number of function evaluations. Differential evolution showed a significantly lower run time for all cases compared to brute force and Bayesian. This is, in part, because it had a stopping criterion that it reached long before it hit its iteration limit, but also because its iteration limit was slightly lower than the other two (due to the nature of DE). Fires 1B and 2B, which doubled the simulation time compared to 1A and 2A, respectively, took over five times more to run due to the exponentially increasing run time of a single fire growth model with increasing simulation time. Similarly, Fires 2A and 2B increased the run time by a factor of four over the 1A and 1B because of the higher wind speed causing higher overall area burned. Higher wind and higher burned area both require more function evaluations internal to the FlamMap software and significantly increase run time.

C. Implications for Wildfire Suppression

The results demonstrate that, for this case, the choice of metric is less critical than the choice of optimizer, but the choice of optimizer is not trivial. Even for the area burned metric, different optimizers find different optimal locations, some quite far from each other. In all cases, a longer simulation time gives a more clear view of the optimal location at the cost of run time. In operational contexts, this tradeoff may be acceptable when simulation time is not a limiting factor, as improved solution quality can translate into more effective suppression outcomes. However, in time-critical scenarios, minimizing function evaluation and reducing simulation time may provide a “good enough” solution at a substantially reduced run time. The variability in optimal drop locations across cases underscores the complexity of wildfire growth and suggests that optimization strategies must be carefully validated before deployment. Ultimately, the choice of optimizer depends on the balance between computational resources, solution quality, and the urgency of decision-making in wildfire response.

VI. Conclusion and Future Work

This study introduced a novel cost function framework for evaluating wildfire suppression strategies, with a focus on optimizing aerial retardant drop locations. By consolidating multiple objectives—including area burned, fire progression, population exposure, and emissions—into a single scalar measure, the framework enables systematic comparison of candidate strategies and provides a unified basis for optimization.

Three gradient-free optimization methods were applied: brute force search, Bayesian optimization, and differential evolution. The results did not show a clear winner for all cases between the three methods. In the majority of cases, Bayesian offered a lower metric value than brute force and differential evolution, but sometimes it required fewer function evaluations and sometimes it required more. Differential evolution was the least robust of the optimizers. In 3 out of 16 cases, it found the optimal solution in fewer function evaluations, but for many cases it reached its stopping criterion early and gave the worst solution of the three. Brute force was in the middle. It improved from first evaluation to last, and for 5 out of 16 cases it found the optimal solution, but without a clear stopping criterion it has very limited run time. Of the three, for these particular simulated fires, Bayesian optimization showed the highest potential to find the optimal drop location if combined with a quality stopping criterion. The variability in optimal drop locations across different cases highlights the complexity of the wildfire growth problem and suggests that multiple local optima may exist within the solution space.

Future work will extend this analysis to a wider range of fire scenarios, explore hybrid optimization approaches, and investigate the stability of solutions across varying environmental conditions. Among those fire scenarios will be higher population maps and higher population scale factors. The one-dimensional design space can be extended to three dimensions, where instead of limiting the drop to a single curve, the x-coordinate, y-coordinate, and angle of the drop can be optimized to find the best drop anywhere on the map. With the expansion of the design space, constraints can be added to make this a constrained optimization problem. Possible constraints include excluding drop locations that are unsafe for a particular aircraft to drop, such as in valleys or up hills, excluding locations too near populated areas, or excluding areas near ground units. Expanding the design space will require a significant speedup of the underlying FlamMap software which is already in progress. Additionally, the optimization can be expanded to not just one drop, but sequences of drops by different types of aircraft. Ultimately, the integration of multi-objective cost functions with advanced optimization techniques has the potential to enhance both the tactical effectiveness and strategic sustainability of wildfire response.

References

- [1] Sullivan, A. L., “A review of wildland fire spread modelling, 1990-present 3: Mathematical analogues and simulation models,” *International Journal of Wildland Fire*, Vol. 18, No. 4, 2009, p. 387. <https://doi.org/10.1071/WF06144>.
- [2] Scott, J. H., “Introduction to Fire Behavior Modeling,” *National Interagency Fuels, Fire, & Vegetation Technology Transfer*, 2012. URL https://pyrologix.com/wp-content/uploads/2014/04/Scott_2012.pdf.
- [3] Andrews, P. L., “The Rothermel surface fire spread model and associated developments: A comprehensive explanation,” *U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station*, 2018. <https://doi.org/10.2737/RMRS-GTR-371>.
- [4] Heinsch, F. A., “Fuel Model,” *Encyclopedia of Wildfires and Wildland-Urban Interface (WUI) Fires*, edited by S. L. Manzello, Springer International Publishing, Cham, 2019, pp. 1–19. URL https://doi.org/10.1007/978-3-319-51727-8_178-1.

- [5] National Interagency Fire Center, “Appendix B: Fire Behavior Supplement,” *NWCG Fireline Handbook*, 2006. URL <https://gacc.nifc.gov/oncc/docs/appendixB.pdf>.
- [6] Richards, G. D., “An elliptical growth model of forest fire fronts and its numerical solution,” *International Journal for Numerical Methods in Engineering*, Vol. 30, No. 6, 1990, pp. 1163–1179. <https://doi.org/10.1002/nme.1620300606>.
- [7] Richards, G. D., “A General Mathematical Framework for Modeling Two-Dimensional Wildland Fire Spread,” *International Journal of Wildland Fire*, Vol. 5, No. 2, 1995, pp. 63–72. <https://doi.org/10.1071/wf9950063>.
- [8] Federal Emergency Management Agency, “What is the WUI?” , Jun. 2022. URL <https://www.usfa.fema.gov/wui/what-is-the-wui/>.
- [9] US Forest Service, “Safe and Effective Wildfire Response,” , Dec. 2016. URL <https://www.fs.usda.gov/managing-land/fire/response>.
- [10] Miller, C., “When It Comes to Wildfire Solutions, Relocating Communities Is a Tough Sell | KQED,” , Aug. 2019. URL <https://www.kqed.org/science/1945874/when-it-comes-to-wildfire-solutions-relocating-communities-is-a-tough-sell>.
- [11] Bacon, A., “Will California homeowners relocate or rebuild? Both are costly,” , Jan. 2025. URL <https://www.accuweather.com/en/weather-news/will-california-homeowners-relocate-or-rebuild-both-are-costly/1733155>.
- [12] North American Moving Services, “How Are Moving Costs Calculated? | Factors Impacting Your Price,” , 2025. URL <https://www.northamerican.com/moving-resources/how-do-moving-companies-charge>.
- [13] Defense, F. W., “Aerial Firefighters & Fire Fighting - Much Needed Support,” , Oct. 2016. URL <https://www.frontlinewildfire.com/wildfire-news-and-resources/aerial-wildfire-fighting-how-effective-is-it/>.
- [14] Rothermel, R. C., “A mathematical model for predicting fire spread in wildland fuels,” *Res. Pap. INT-115. Ogden, UT: U.S. Department of Agriculture, Intermountain Forest and Range Experiment Station. 40 p.*, Vol. 115, 1972. URL <https://research.fs.usda.gov/treesearch/32533>.
- [15] “WindNinja,” 2024. URL <https://research.fs.usda.gov/firelab/products/dataandtools/windninja>.
- [16] Doerr, S. H., and Santín, C., “Global trends in wildfire and its impacts: perceptions versus realities in a changing world,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, Vol. 371, No. 1696, 2016, p. 20150345. <https://doi.org/10.1098/rstb.2015.0345>.
- [17] XPrize Wildfire, “State of Prize Report: The Grand Challenge and Rationale beh,” Tech. rep., Jun. 2023. URL <https://www.xprize.org/prizes/wildfire/articles/state-of-prize-report-the-grand-challenge-and-rationale-behind-xprize-wildfire>.
- [18] National Wildfire Coordinating Group, “Wildland Fire Suppression Tactics Reference Guide,” , Apr. 1996.
- [19] Finney, M. A., “FARSITE: Fire Area Simulator-model development and evaluation,” *Res. Pap. RMRS-RP-4, Revised 2004. Ogden, UT: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station. 47 p.*, Vol. 4, 1998. <https://doi.org/10.2737/RMRS-RP-4>.
- [20] Finney, M. A., “Fire growth using minimum travel time methods,” *Canadian Journal of Forest Research*, Vol. 32, No. 8, 2002, pp. 1420–1424. <https://doi.org/10.1139/x02-068>.
- [21] Finney, M. A., “An Overview of FlamMap Fire Modeling Capabilities,” 2006. URL <https://www.semanticscholar.org/paper/An-Overview-of-FlamMap-Fire-Modeling-Capabilities-Finney/a61989f226731b2d1403eb4356c6eea652e27698>.
- [22] “FlamMap,” 2024. URL <https://research.fs.usda.gov/firelab/products/dataandtools/flammap>.
- [23] Martins, J. R. R. A., and Ning, A., *Engineering Design Optimization*, 2nd ed., Cambridge University Press, Cambridge, UK, 2021. URL <https://mdobook.github.io>.
- [24] Gabbert, B., “Happy Camp Fire Complex closing in on "megafire" status,” , Sep. 2014. URL <https://wildfiretoday.com/2014/09/05/happy-camp-fire-complex-closing-in-on-megafire-status/>.
- [25] Cal Fire, “Happy Camp Complex,” , Dec. 2023. URL <https://www.fire.ca.gov/incidents/2023/8/16/happy-camp-complex>.
- [26] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, Cambridge University Press, Cambridge, UK ; New York, 2007.

- [27] Scott, J. H., and Burgan, R. E., "Standard fire behavior fuel models: a comprehensive set for use with Rothermel's surface fire spread model," Tech. Rep. RMRS-GTR-153, U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station, Ft. Collins, CO, 2005. <https://doi.org/10.2737/RMRS-GTR-153>.
- [28] European Commission (ed.), *GHSL data package 2023*, Publications Office, Luxembourg, 2023. <https://doi.org/10.2760/20212>.
- [29] Missoula Fire Sciences Laboratory, "FOFEM/SpatialFOFEM - fire effects model," , Aug. 2020. URL <https://research.fs.usda.gov/firelab/products/dataandtools/fofem/spatialfofem-fire-effects-model>.