

Article

Investigating the Impact of Aerial Firefighting on Rate of Wildfire Spread

Lindsay Wiard-Greene ^{1,*}, Jesse Johnson ¹, John Hogland ², Fredrick Bunt ¹ and Jake Bova ¹¹ Department of Computer Science, University of Montana, Missoula, MT 59812, USA² Rocky Mountain Research Station, U.S. Forest Service, Missoula, MT 59801, USA; john.s.hogland@usda.gov

* Correspondence: lindsay.wiard@umontana.edu

Abstract

Aerial retardant drops are widely used in wildfire suppression, yet their effectiveness in slowing fire spread remains difficult to quantify at scale. This study evaluates their impact on wildfire rate of spread (ROS) using a framework that combines observed and counterfactual (synthetic) drop locations from 62 Oregon wildfires. Synthetic drops were generated to simulate a no-suppression baseline, enabling comparison of ROS with and without suppression. We trained two random forest classifiers: one with real and synthetic drops (full model) and one with only synthetic drops. Both incorporated environmental and topographic features to predict whether spread slowed following a drop. While the full model performed well, the real-synthetic indicator had low feature importance, offering limited causal evidence that aerial suppression consistently reduced spread. The synthetic-only model produced similar performance, suggesting that drops with observed ROS reductions often coincided with favorable environmental and topographic conditions and may have occurred independent of suppression. These findings highlight the challenges of evaluating suppression at scale and emphasize the need for finer data, detailed operational records, and advanced modeling to better assess the role of aerial fire retardant drops in future wildfire management activities.

Keywords: suppression; aerial retardant drops; fire management; rate of spread; causal inference; counterfactual modeling; machine learning

1. Introduction

Over the past few decades, wildfires in the United States have increased in size, frequency, and intensity, presenting growing challenges for fire management and suppression efforts. Research indicates that the “most extreme fires are also larger, more common, and are more likely to co-occur with other extreme fires”, further exacerbating the difficulty of wildfire suppression [1]. This trend is projected to worsen, with future wildfire seasons expected to be longer and characterized by more frequent and severe fire events [2].

As wildfires continue to grow in scale, suppression efforts have increasingly relied on aviation support, including aerial retardant drops and water scooping operations. These aerial suppression efforts come with high costs—monetary, environmental, and human. Between 2007 and 2011, the U.S. Forest Service (USFS) spent approximately \$323 million annually on aviation-based suppression efforts [3]. This expenditure increased in the following decade, with aviation costs averaging \$493 million per year between 2011 and 2020, accounting for approximately 30% of the total USFS wildfire suppression expenditures [4]. These high suppression costs often divert funding away from proactive fire



Academic Editor: W. John Braun

Received: 5 November 2025

Revised: 17 December 2025

Accepted: 18 December 2025

Published: 19 December 2025

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

prevention strategies, such as prescribed burns, forest thinning, and community resilience programs—measures that are generally more cost-effective and reduce long-term risk [5].

Beyond financial costs, aerial suppression also has significant personnel-related consequences and potential environmental impacts. From 2011 through 2020, aviation-related accidents accounted for the highest percentage of federal firefighter fatalities, including both contract and employee personnel, representing 30% of all fatalities during that period [4]. Furthermore, the use of fire retardant has introduced substantial amounts of toxic metals into ecosystems. It is estimated that from 2009 to 2021, aerial suppression operations contributed approximately 380,000 kg of toxic metals to the environment [6].

In addition to these high costs and risks, it must be considered that aviation resources are also limited, with a finite number of aircraft and personnel available. Given these challenges, it is critical to assess the effectiveness of aerial suppression efforts. Understanding when and where aviation resources have the most impact can help optimize their use, ensuring that they provide maximum benefit while minimizing unnecessary expenditures, environmental degradation, and personnel hazards. This study aims to evaluate the effectiveness of aviation in wildfire suppression and identify conditions under which it has the greatest operational value.

Considering the critical importance of assessing the effectiveness of aerial firefighting, researchers have employed various approaches to achieve this goal. In his review of wildfire suppression effectiveness research, Plucinski [7] explores the research completed prior to 2016. He notes that previous studies have been conducted at multiple scales, ranging from flame-scale experiments to fireline and fire incident-scale analyses. Flame-scale studies, examining areas less than 10 m in length, have typically focused on retardant effectiveness in terms of flame length or intensity. These experiments are easier to design and control, allowing researchers to isolate specific factors influencing suppression outcomes. At the fireline scale (approximately 10–1000 m in length), researchers have investigated suppression effects along sections of the fire perimeter to evaluate resource productivity and the impact of tactics on fire behavior in the field. Fire studies at the incident-scale, which refers to the entirety of the burned area, have further expanded insight into aerial support effectiveness by analyzing broader suppression efforts, including initial attack containment success.

Due to the nature of the data available for this study, our research focuses on fireline and fire incident scales. At these scales, controlled experiments are challenging to conduct, necessitating reliance on observational studies. Therefore, careful study design is necessary to account for confounding variables and ensure meaningful conclusions. Researchers have proposed various methodological approaches to improve the reliability of suppression effectiveness assessments. One approach to examining suppression effects on fire behavior is by comparing fire behavior along suppressed sections of the fire perimeter to unsuppressed sections, provided that comparable conditions—such as fuel, terrain, and weather—can be established [7].

In their paper, Plucinski and Pastor [8] state that to attribute fire behavior changes to suppression efforts, it is preferable to account for other influencing variables such as aircraft flight and delivery system conditions, vegetation and fuel characteristics, weather and terrain conditions, fire behavior metrics, and concurrent ground suppression activities. Although many of these variables undoubtedly influence fire behavior in reality, they may not contribute meaningfully to improving analysis or predictive performance in a modeling context. Additionally, even with all these factors considered, fully attributing suppression effectiveness to aerial efforts remains challenging. Therefore, we should aim to collect as many variables as possible and prioritize those that are most relevant for improving our analysis.

Many prior fire suppression studies have relied on in situ data collection, where researchers gather detailed information about individual fire incidents. This approach allows for a more detailed analysis but often limits the number of cases that can be examined, making it difficult to draw broad conclusions about suppression effectiveness. Despite these limitations, several studies have leveraged in situ data to evaluate fire suppression tactics.

The U.S. Forest Service attempted to expand the available data by completing the Aerial Firefighting Use and Effectiveness (AFUE) [9] study which provided empirical data on the effectiveness of aerial suppression efforts. From 2015 to 2018, AFUE personnel collected data across all nine Forest Service regions, documenting drop objectives and outcomes, aircraft types, fire characteristics, and ground resource involvement. While the AFUE report summarized statistics on the performance of drops, it did not analyze the conditions under which those drops were most effective, limiting its utility in determining optimal deployment strategies.

Katuwal et al. [10] aimed to provide broad recommendations on when different fire suppression tactics, including aviation, might be most effective by using stochastic frontier analysis to evaluate the efficiency of various tactics. The study assessed factors contributing to the maximization of controlled fireline length. Analyzing daily perimeters from 63 fires, the study examined the impact of crews, dozers, engines, helicopters, and air tankers. The results indicated that air resources (helicopters and air tankers) did not significantly contribute to controlled fireline length. On the other hand, the authors noted that the study could not account for the specific mission objectives of these aircraft, which may have aimed to slow fire progression rather than halt it entirely. This limitation underscores the challenge of evaluating suppression effectiveness when specific operational goals are unknown.

In a study out of Ontario, Wheatley et al. [11] considered wildfires that received aviation support and modeled initial attack (IA) success using logistic regression with supervised forward-selection of variables. IA success was defined as a fire being contained by 1:00 PM on the second day of attack. This analysis identified factors associated with containment success. While volume of retardant dropped and the number of airtankers used were negatively associated with suppression success (likely due to more challenging fires receiving greater resources), other key predicting variables included fire size at IA initiation, weather conditions, fire behavior, fuel type, and fire cause. The study suggested that redirecting IA air resources toward fires with their model's highest containment probability could improve suppression efficiency.

While in situ data has been invaluable, recent advancements in remote sensing have enabled researchers to analyze suppression effectiveness at broader scales, reducing data collection costs and increasing the number of incidents studied. The next couple of studies have leveraged remotely sensed data to further refine suppression assessments, offering insights that were previously difficult to obtain with ground-based observations alone.

Stonesifer et al. [12] analyzed historical large airtanker drop data from the 2010–2012 fire seasons across the United States by geospatially intersecting drop locations with operational and environmental factors. Consistent with previous research, results indicate that aviation resources are predominantly deployed on fires that are more difficult to contain. Airtankers were generally used according to operational guidelines, including appropriate altitude and speed. However, findings suggest that drops sometimes occurred under conditions where flame lengths exceeded recommended thresholds. The study also highlights the influence of human proximity on drop locations and reveals that aviation resources were utilized during extended attack operations more frequently than previously recognized.

In his PhD dissertation for Colorado State University, Bryan [13] employed a difference-in-differences estimation and a regression discontinuity design to compare fire behavior in treated versus untreated areas. His findings suggest that large air tankers effectively

diminish flame intensity, reduce wildfire spread likelihood, and delay fire growth. He also found that aerial retardant drops were most effective in grass and shrub fuel types and when used in conjunction with other suppression efforts. Furthermore, Bryan's research introduced a cost-effectiveness metric that considers avoided damages along with providing a framework for evaluating the trade-offs between the usage of aerial and ground firefighting resources.

The next step in suppression research is the development of strategic, risk-informed decision support frameworks that integrate insights from both in situ and remotely sensed studies. Some of these frameworks already exist, such as the Aviation Use Summary (AUS), which was reviewed in Stonesifer et al. [4]. The AUS leverages aircraft event tracking data, geospatial information and analytical tools to evaluate aircraft use at the incident level, helping guide risk based decision making.

The ultimate goal of this study is to develop a predictive model that can help determine the conditions under which an aerial retardant drop may reduce the rate of spread (ROS) of a wildfire. Identifying these conditions would improve strategic deployment of aviation resources, optimizing their impact while mitigating unnecessary costs and risks. Our research aligns with earlier studies by utilizing topographical, weather, fuel, and fire behavior variables to inform aerial suppression decision-making frameworks. While only a few recent studies have capitalized on the availability of remotely sensed data, our approach aims to build upon this foundation. Our research is innovative because it has access to extensive new aviation data, will calculate the ROS using progressive 12-h perimeters, and will use ROS to evaluate the effectiveness of aerial interventions.

The remainder of this paper is organized as follows. In Section 2, we describe the data sources, preprocessing steps, and modeling framework used to evaluate the effectiveness of aerial retardant drops in reducing wildfire ROS. This includes our approach for calculating ROS, generating synthetic drops, and building predictive models. Section 3 presents the performance of the models we developed to assess the impact of aerial suppression on ROS reduction. In Section 4, we interpret these findings in the context of wildfire suppression effectiveness, highlight key limitations of our approach, and compare our results to those of prior studies. Finally, in Section 5, we outline directions for future research and the potential implications of our findings for operational fire management.

2. Materials and Methods

2.1. Data Collection

To evaluate the effectiveness of aerial retardant drops, we assembled a comprehensive dataset integrating multiple sources of fire behavior and suppression data. Data were collected in support of our effort to distinguish predictable fire behavior from the fire's response to retardant. First, we obtained access to a dataset collected by the U.S. Forest Service detailing large airtanker drop locations and characteristics across the United States. This dataset provided critical information on drop timing, location, and operational parameters. Additionally, our lab had previously compiled a dataset containing topographical, fuel, and weather information for wildfires in Oregon, allowing us to analyze environmental factors influencing suppression outcomes. To further assess the impact of aerial suppression on fire behavior, we calculated ROS by acquiring progressive 12-h fire perimeter data accessed through NASA's Earthdata portal.

The dataset detailing large airtanker drops was collected across the United States by both federally owned and contracted airtankers. Drop events were recorded using Additional Telemetry Unit (ATU) sensors, which automatically logged the location of door openings and closings associated with retardant releases. In addition to drop locations, ATUs captured key flight and delivery parameters, including airspeed, altitude, heading,

and the volume of retardant delivered. These door event data points allowed for the creation of precise drop line representations within a geographic information system (GIS). While ATU data collection is fully automated, the dataset we used had been previously refined by Stonesifer. She had removed false positives, manually linked drop events to specific wildfire incidents, and flagged potential non-fire drops. It contained data for approximately 7500 drops spanning the nation from 2016 through 2021.

Prior to this study, our lab had compiled a comprehensive dataset encompassing environmental variables for fire-affected regions. This dataset was created by extracting topographical, weather, and fuel characteristics for points within wildfires in Oregon. The wildfires were identified using the Monitoring Trends in Burn Severity (MTBS) fire perimeters dataset from 2016 to 2020. The MTBS program is an interagency initiative aimed at consistently mapping the burn severity and extent of large fires (greater than 1000 acres in the western United States) [14]. This dataset provided critical contextual information for understanding the conditions under which aerial suppression efforts were deployed. The specific variables included in this dataset, along with their sources and spatial resolutions, are detailed in Table 1.

Table 1. Summary of environmental features, including topography, weather, and fuels data.

Feature	Resolution	Source
Topography		
Elevation	30 m, Constant	EDNA [15]
Slope	30 m, Constant	EDNA [15]
Aspect	30 m, Constant	EDNA [15]
Flow Accumulation	30 m, Constant	EDNA [15]
Hillshade	30 m, Constant	Derived from Slope, Aspect (zenith = 45°, azimuth = 180°)
Weather		
Solar Radiation	4 km, weekly average	gridMET [16]
Vapor Pressure Deficit	4 km, weekly average	gridMET [16]
Palmer Drought Severity Index	4 km, weekly average	gridMET [16]
Mean Temperature (Warmest Month)	1 km, monthly average	AdaptWest [17]
Mean Temperature (Coldest Month)	1 km, monthly average	AdaptWest [17]
Temperature Difference (Coldest vs. Warmest Month)	1 km, difference of monthly averages	AdaptWest [17]
Mean Annual Temperature	1 km, yearly average	AdaptWest [17]
Maximum Temperature	1 km, monthly max	Daymet: Monthly [18]
Minimum Temperature	1 km, monthly min	Daymet: Monthly [18]
Fuels		
LANDFIRE Vegetation Type	30 m, 2014 Update	LANDFIRE [19]
LANDFIRE Fuels Model 40	30 m, 2014 Update	LANDFIRE [19]
Perennial Herbaceous Biomass	30 m, yearly average	Rangeland Analysis Platform [20]
Annual Herbaceous Biomass	30 m, yearly average	Rangeland Analysis Platform [20]
Normalized Difference Vegetation Index (NDVI)	4 km, weekly average	NOAA CDR Program [21]

To track fire progression and quantify ROS, we used fire perimeter data generated by the Fire Events Data Suite (FEDS) algorithm, accessed through NASA's Earthdata portal. The FEDS algorithm processes thermal observations from the Visible Infrared Imaging Radiometer Suite (VIIRS) sensors aboard the Suomi NPP and NOAA-20 satellites

to delineate fire perimeters. With twice-daily overpasses, VIIRS provides updated fire perimeter data every 12 h, offering a detailed view of fire evolution over time [22]. We calculated the ROS by rasterizing each fire perimeter at a 30-m resolution and sequentially dilating the perimeter until it intersects the subsequent perimeter. The number of dilations required to reach each pixel in the subsequent perimeter was multiplied by 30 m and divided by the time difference in hours between observations, yielding a ROS (m/h) for each point along all perimeters, excluding the first. As an illustrative example, Figure 1a shows the progression of 12-h fire perimeters for the Green Ridge Fire (August 2020, Oregon), and Figure 1b shows the resulting ROS values calculated along the eighth perimeter of that fire.




Figure 1. The Green Ridge Fire, which occurred in August 2020 in Oregon. (a) A typical progression of 12-h perimeters. The shading from white to dark red represents the fire's spread over time, with white indicating the initial perimeter and dark red representing the final perimeter. (b) The extent of the seventh perimeter, shaded in white, with rate of spread (ROS) values (m/h) calculated at each point along the following perimeter.

2.2. Data Processing

To integrate the three datasets (the drops, the features, and the VIIRS perimeters), we first limited our analysis data within Oregon between 2018 and 2020 for which all required data sources were available. We resolved differences in naming conventions among the datasets by identifying matching fire events and drops based on location and timing. We then associated each airtanker drop with the corresponding environmental and fire behavior data. Environmental variables were extracted from the precompiled feature dataset (Table 1) by buffering each drop location by 20 m and averaging the values of all intersecting raster cells.

We then quantified fire behavior before and after each airtanker drop using ROS derived from VIIRS fire perimeters delineated at a 12-h temporal resolution. For each drop, we identified the fire perimeter immediately preceding and following the drop based on time. For each point along the drop line, we located the closest points on these perimeters and averaged the ROS from those locations. These ROS values therefore represent the 12-h spread rate from one perimeter to the next. If a drop occurred after the final recorded perimeter, we assigned a post-drop ROS of zero. We excluded drops where both pre- and post-drop ROS were zero, as these cases were unlikely to provide meaningful information on suppression effectiveness. Finally, we quantified the change in ROS by subtracting the pre-drop ROS from the post-drop ROS, where negative values indicate a reduction in fire spread following the drop. The final dataset includes 586 drops associated with 31 fires.

We began our analysis by examining the overall distribution of ROS changes for all recorded drops. To visualize this, we present a histogram of the change in ROS in Figure 2. The average change in ROS across all drops was -70 m/h, indicating a general reduction in fire spread following aerial suppression. However, a significant number of drops resulted in little to no change in ROS. Given this, we established a threshold to identify drops

that effectively slowed fire spread. Specifically, we consider only those drops where ROS decreased by at least 70 m/h as successful in slowing fire progression. We refer to cases having less than this threshold of change of ROS as persisting.




Figure 2. Histogram of change in ROS following aerial retardant drops. The vertical dashed red line represents the mean change in ROS (-70 m/h). Drops to the left of this threshold are considered effective in slowing fire spread. Drops to the right are considered to have persistent ROS.

Next, we compared the distributions of the environmental variables for drops that effectively slowed the ROS to those where the ROS persisted. Our analysis revealed notable differences in several variables, including temperature difference and elevation. As can be seen in Figure 3a,b, these variables exhibited distinct distributions between the two groups as can be seen in the means and standard deviation. This suggested that they may serve as important predictors in explaining variation in suppression outcomes. In contrast, other variables, such as the previous area of the fire and hillshade, exhibited relatively similar distributions (Figure 3c,d). Given the potential for machine learning models to detect subtle interactions and patterns not immediately evident through visual inspection, we retained all variables in our modeling framework. The inclusion of a broader set of features allowed the models to explore complex relationships within the data.

To rigorously assess whether aerial retardant drops contribute to reductions in the ROS, it is essential to establish whether fire behavior would have remained unchanged in the absence of a drop. Without a clear comparison or counterfactual case, attributing changes in ROS to suppression efforts remains speculative. We address this limitation by introducing synthetic drops that provide a baseline for expected fire behavior without aerial suppression. By comparing changes in ROS between real and synthetic drops, we can better discern whether observed reductions in ROS exceed those occurring naturally in unsuppressed areas. This approach enables us to better quantify the contribution of aerial suppression to changes in fire spread.

The synthetic drops were designed to mirror the temporal distribution, frequency, and spatial scale of observed drops while remaining independent of fire behavior outcomes. We first identified all fires with documented aerial retardant drops ($n = 31$) and all fires without drops ($n = 34$). Fires in both groups were characterized by the number of available progressive fire perimeters, which served as a proxy for fire duration. Fires without drops were randomly subsampled and matched to fires with drops based on the number of perimeters, ensuring comparable fire lengths between fires with documented and synthetic drops.




Figure 3. Probability density distributions of environmental variables for aerial drops. The blue distributions represent drops with a slowed ROS, while the orange distributions indicate drops with a persistent ROS. The thick vertical lines represent the distribution means, while the thin dashed lines denote one standard deviation away from the mean.

For each fire with observed drops, we quantified the distribution of drops across its progressive perimeters. Specifically, the perimeter index on which each drop occurred was converted to a proportional fire progression metric (e.g., a drop occurring on perimeter 2 of 5 corresponds to 40% of fire progression). The number of drops occurring at each proportional perimeter position was then transferred to the matched fire without drops. For example, if a fire with five perimeters contained three drops at 40% progression and six drops at 60% progression, the matched control fire received three synthetic drops on its 40% perimeter and six synthetic drops on its 60% perimeter. An example of a fire with observed drops and its matched fire with synthetic drops is shown in Figure 4.

Synthetic drop lengths were generated by sampling from an empirical probability distribution constructed from the lengths of observed retardant drops. For each synthetic drop, candidate lengths were restricted to values less than one-tenth of the diagonal length of the target perimeter's bounding box to prevent unrealistically long drops relative to perimeter scale. A length was randomly selected from the constrained distribution and removed from the candidate pool to avoid reuse.

The spatial placement of synthetic drops along the fire perimeters was executed using a randomized approach subject to several constraints. Each target perimeter was buffered by a fixed distance, and a random starting point was selected within the buffered region. An endpoint was then generated at the sampled drop length distance from the starting point, with both endpoints required to fall within the buffer. Synthetic drops were explicitly prohibited from being placed in locations where both pre- and post-drop ROS values were zero, as such areas were excluded from the original set of real drops and would not provide

a comparable basis for analysis. Additionally, each synthetic drop was required to intersect areas containing valid environmental and topographic covariate data to ensure consistency with subsequent feature extraction.




Figure 4. Comparison of observed and synthetic aerial retardant drop placement on matched fires. On the (left), the Tenino Fire is shown with observed aerial retardant drops, and on the (right) the matched unnamed fire is shown with synthetic drops. In both panels, a typical progression of 12-h fire perimeters is displayed, with shading from white to dark red representing fire spread over time. Both fires contain five progressive perimeters, with 11 drops associated with the third perimeter, five with the fourth, and one with the fifth. The synthetic drops preserve the temporal distribution of real drops while being randomly placed along corresponding perimeters subject to spatial and environmental constraints.

Following placement, the same feature extraction and processing workflow applied to real drops was applied to all synthetic drops. This procedure resulted in a total of 586 synthetic drops distributed across 31 fires, matching the number of real drops included in the analysis. The distribution of changes in ROS for synthetic drops closely resembled that of the real drops (Figure 5). To formally assess this similarity, we conducted a Welch's unequal-variance *t*-test comparing changes in ROS between real and synthetic drops. The resulting *p*-value (0.5046) indicated no statistically significant difference between the means, suggesting that the synthetic drops effectively captured the overall fire behavior characteristics of the observed drops.

2.3. Model Creation

To assess the effectiveness of aerial retardant drops in slowing fire spread, we developed predictive models that classified whether a drop effectively reduced the ROS. Effectively reducing ROS was measured by decreasing the ROS by at least 70 m/h between consecutive perimeters. A key component of one of our modeling approaches was incorporating both real and synthetic drops in the analysis. Specifically, we hypothesized that if the feature indicating whether a drop was real or synthetic emerged as a major predictor, it would suggest that actual retardant drops had a significant impact on ROS reduction beyond what would be expected from natural fire behavior alone.

To implement this model, which we will refer to as the full model, we utilized a random forest classifier due to its built-in feature importance metrics, robustness to non-linear relationships, and ability to handle interactions between variables. The model was trained using the environmental features from Table 1, along with three additional predictor variables being the length of the drop, the area of the fire within the perimeter preceding the

drop, and an indicator variable distinguishing real drops from synthetic ones. The target variable was a binary classification of whether or not a drop effectively slowed the ROS. To ensure a fair evaluation, we included both real and synthetic drops in the training and testing datasets. We carefully split the data so that all drops associated with a given fire were entirely within either the training or testing set, preventing data leakage and ensuring that the model was tested on independent fires it had not encountered during training.




Figure 5. Histogram of change in ROS. The blue distribution represents the real drops, while the orange distribution indicates the change in ROS for the synthetic drops. To aid in visualization, the change in ROS was restricted to within three standard deviations from the mean.

Since our primary interest was in evaluating the relative importance of the real versus synthetic drop indicator, we first trained a random forest model using all available predictors, then performed an additional feature selection step to reduce redundancy and improve interpretability. Using Spearman rank-order correlation, we quantified pairwise correlations among all predictor variables and applied hierarchical clustering with Ward’s linkage to group highly correlated features. From each cluster, we selected a representative variable based on its impurity-based feature importance in the initial model along with its distributional differences between drops with slowed and persistent ROS (Figure 3). This process yielded a final set of four features: temperature difference, elevation, previous fire area, and hillshade. Additionally, the indicator for real versus synthetic drops was retained regardless of correlation structure, given its central role in the analysis.

As well as training the model on a mix of real and synthetic drops, we also explored an alternative approach where the model was trained exclusively on synthetic drops using all available environmental, topographic, and contextual predictor. The rationale behind this synthetic-only model was to establish a baseline that captured patterns of fire behavior in the absence of aerial suppression. We then applied the trained model to real drops to estimate where ROS would likely have persisted without aerial intervention. This allowed us to infer how frequently retardant drops were effective in slowing ROS, offering a broader perspective on suppression outcomes.

2.4. Model Sensitivity Testing

For both modeling approaches, the full and synthetic-only, we fine-tuned the hyperparameters using a grid search strategy to optimize performance. There were significantly more drops experiencing persistent ROS than those that successfully slowed it which created a large class imbalance in the target variable. Due to this imbalance, the model tended to overpredict the majority class—persistent ROS. To mitigate this, we used balanced accuracy as our scoring metric during grid search. Balanced accuracy accounts for both

sensitivity (true positive rate) and specificity (true negative rate) by averaging recall across classes, making it more appropriate than standard accuracy in imbalanced classification tasks. The equation used to compute this metric is shown in Equation (1). This approach helped us select hyperparameters that more fairly represented both classes and improved our model's ability to more accurately classify drops.

$$\text{Balanced Accuracy} = \frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \quad (1)$$

Specifically, we tuned the number of trees in the forest, the maximum depth of each tree, the minimum number of samples required to split an internal node, the minimum number of samples required to be at a leaf node, along with class weights. The selected hyperparameters for each model are summarized in Table 2.

Table 2. Optimal hyperparameters for the random forest models.

Hyperparameter	Full Model	Synthetic-Only Model
Number of trees	100	10
Maximum tree depth	2	10
Minimum number of samples to split a node	2	2
Minimum number of samples at leaf node	1	8
Class weight	balanced *	balanced *

* The 'balanced' option adjusts weights inversely proportional to class frequencies: $weight_i = \frac{n_{\text{samples}}}{n_{\text{classes}} \cdot n_i}$, where n_i = the number of samples in class i .

We acknowledge that selecting the mean change in ROS as the threshold value (−70 m/h) to distinguish between slowed and persistent ROS is somewhat arbitrary. A change in ROS of 0 m/h may seem like an intuitive dividing line, and we tested several thresholds near zero. However, these near-zero thresholds produced models with little predictive power. This is likely because small variations in ROS, exacerbated by the approximations and uncertainties in our data and calculations, were often not meaningful. As a result, we opted to allow for some margin around this cutoff, choosing a more conservative threshold that better captured substantial reductions in fire spread.

Specifically, we evaluated a symmetric range of threshold values around the mean, varying in 0.05 standard deviation increments up to ± 0.25 . For both models, performance was fairly stable and consistent within ± 0.1 standard deviations, with classification metrics remaining high and balanced. As thresholds deviated further, performance began to decline moderately, though it remained within an acceptable range. Even so, when we tested the median as a threshold, which is more than 0.3 standard deviations above the mean (approximately −12 m/h), performance dropped significantly. This confirmed that thresholds that were too close to zero, such as the median, reduced the model's discriminative power by classifying many marginal changes as successful.

In addition to adjusting the threshold within a binary framework, we also explored a multiclass classification approach. Drops were grouped into classes such as slowed, constant, and persistent ROS, based on the magnitude of ROS reduction. However, the distribution of values was heavily concentrated near zero, resulting in a large number of drops being assigned to the constant category. This class imbalance led to a model that achieved high overall accuracy but poor balanced accuracy, as it predominantly predicted that all drops would fall into the majority class. Attempts to balance the class distribution reduced overall accuracy and failed to meaningfully improve performance, likely because overlapping feature patterns across classes made consistent differentiation difficult.

We also tested regression models, including both linear regression and random forest regression, to directly predict the change in ROS. Both approaches yielded low predictive power, with R^2 values below 0.1. This poor performance was again likely due to the high concentration of ROS change values near zero, limiting the model's ability to distinguish meaningful patterns in the continuous outcome.

Taken together, these results suggest that while the mean-based threshold for the binary classification is not optimal, it lies near a local optimum in terms of model performance. Thresholds set too far from the mean, particularly in the more lenient direction, compromised predictive performance. Thresholds placed too far from zero are difficult to justify conceptually and lead to an even greater imbalance between classes, further limiting model reliability. Ultimately, the binary classification using the -70 m/h threshold offered the best balance between interpretability and model performance.

3. Results

3.1. Full Model and Feature Importance

First, we summarize the results of the full model, which was trained on both real and synthetic drops and a limited set of predictor features: temperature difference, elevation, previous fire area, hillshade, and the real vs. synthetic drop indicator.

Accuracy: 0.8802 Balanced Accuracy: 0.8179

The full model attained an accuracy score of 0.8802, indicating that it correctly classified whether a drop effectively reduced the ROS in approximately 88% of cases. Because of the imbalance in target variable classes, we also examined the balanced accuracy. The balanced accuracy for this model was 0.8179, suggesting that while the model performed well overall, it had slightly lower effectiveness when accounting for class imbalance. The confusion matrix found in Table 3 provides further insight into the model's classification performance. The upper left and lower right corners represent accurate predictions, while the lower left and upper right corners represent misclassified cases. Overall, the model correctly identified 74% of slowed cases and 90% of persistent cases, with misclassification rates of 26% and 10%, respectively. This matrix indicates that the model produced a significant number of correct predictions, with only a small percentage of errors.

Table 3. Full Model Confusion Matrix.

	Actual Slowed	Actual Persistent
Predicted Slowed	31	35
Predicted Persistent	11	307
Total	42	342

To understand the relative contribution of each predictor, we first examined the impurity-based feature importance values from the random forest classifier (Figure 6). This metric quantifies how much each feature reduces impurity, or uncertainty, when making splits in the decision trees. Among the five features used in the model, the variables that showed distinct distributions between drops with slowed and persistent ROS, such as temperature difference and elevation (Figure 3), ranked highest in importance, reinforcing their influence on model predictions. Features with more uniform distributions across outcomes, like drop length and slope, contributed less. Most notably, the indicator distinguishing real from synthetic drops was the least important feature, suggesting limited influence of aerial suppression on ROS outcomes as learned by the model.




Figure 6. Feature importance values from the full model. Each bar represents the relative importance of the predictor in reducing impurity across the decision trees. The real drop indicator is highlighted in orange and is among the least influential features.

To validate and expand upon these findings, we also calculated permutation-based feature importance (Figure 7). Unlike impurity-based methods, permutation importance measures the decrease in model performance when the values of a single feature are randomly shuffled, effectively breaking its relationship with the target variable. This approach provides a more direct estimate of how much each feature contributes to predictive accuracy and is less sensitive to biases related to variable scale or cardinality. The permutation results mirrored the impurity-based ranking exactly, with temperature difference and elevation again emerging as the top contributors and the real-versus-synthetic drop indicator remaining last. This consistency across two independent importance measures strengthens the conclusion that environmental and topographic context, rather than suppression activity, was the primary driver of model predictions.




Figure 7. Permutation-based feature importance values from the full model. Each boxplot shows the distribution of accuracy decrease when a given feature is randomly permuted across multiple iterations. Larger decreases indicate greater importance. Consistent with impurity-based results, the real drop indicator had the lowest impact on model accuracy.

3.2. Synthetic-Only Model and Real Drop Evaluation

Next, we summarize the results of the model trained exclusively on synthetic drops.

Accuracy: 0.8182 Balanced Accuracy: 0.8735

These high accuracy scores indicate that the model effectively learned patterns of fire behavior without the influence of aerial suppression. A confusion matrix provided in Table 4 shows that the model correctly identified 94% of slowed cases and 81% of persistent cases, with misclassification rates of 6% and 19%, respectively. This demonstrates strong predictive performance in the absence of retardant.

Table 4. Synthetic-Only Model Confusion Matrix.

	Actual Slowed	Actual Persistent
Predicted Slowed	16	33
Predicted Persistent	1	137
Total	17	170

To assess the potential impact of aerial retardant drops, we applied this model to real drops. Table 5 summarizes the results. When considering this figure, it is important to note that applying the synthetic-only model to real drops does not provide a measure of the model's accuracy. The model's purpose is to estimate how fire would behave in the absence of aerial suppression. As such, these results should be interpreted as counterfactual estimates rather than as indicators of predictive performance.

Table 5. Synthetic-Only Model Applied to Real Drops.

	Actual Slowed	Actual Persistent	Total
Predicted Slowed	101	97	198
Predicted Persistent	57	331	388
Total	158	428	586

That said, a few descriptive statistics are worth highlighting. Among the 158 drops where the fire's ROS actually slowed, the model predicted that 101 of them (64%) would have a slowed ROS. Additionally, the model predicted that ROS would persist for 388 of the real drops. These locations are particularly noteworthy, as they represent areas where, according to the model, fire spread would likely have continued in the absence of retardant. In practice, only 57 of these 388 drops (15%) were associated with a slowed ROS.

4. Discussion

Despite the strong overall performance of the model trained on both real and synthetic drops, the results do not provide strong evidence that aerial suppression efforts were a primary driver in reducing the ROS over the 12-h periods used for model evaluation. The inclusion of synthetic drops allowed us to assess whether real drops had a distinct impact beyond what would be expected from natural fire behavior. Nonetheless, the indicator variable distinguishing real from synthetic drops was among the least important predictors in the model. This suggests that the use of retardant did not meaningfully influence the model's ability to predict whether ROS would be reduced 12 h later, and that environmental and contextual variables played a far more significant role.

The evaluation using the synthetic-only model further reinforces the lack of clear evidence that aerial suppression efforts significantly contributed to reducing ROS. Among the drops where the fire's ROS actually slowed, the model predicted that 64% would have slowed even in the absence of retardant. This suggests that, in many cases where suppression efforts appeared to be successful, the observed slowing may have occurred regardless of intervention, aligning with conditions the model associated with natural deceleration. Furthermore, if aerial retardant had been a strong driver in reducing fire spread, we would

have expected a substantial number of the predicted-to-persist cases to exhibit slowed ROS in reality. However, only about 15% of those drops actually resulted in reduced ROS, while a large majority continued to exhibit persistent spread. This outcome suggests that in most cases, the presence of a retardant drop did not correspond to a meaningful deviation from the fire behavior expected without suppression.

To further evaluate whether the inclusion of real drops influenced model predictions, we applied the full model and the synthetic-only model to the real drops in the testing set and compared their classification outcomes. **Using McNemar's test, a statistical method designed to assess differences between paired categorical outcomes, we found no significant difference between the two models' predictions ($p = 0.774$).** This result suggests that, despite the full model having access to the real drop indicator variable, its predictions on real drops were statistically indistinguishable from those of the synthetic-only model, further reinforcing the limited predictive value of the real drop flag.

Overall, the findings do not provide strong empirical support for the conclusion that aerial suppression efforts were a consistent or substantial factor in reducing rate of spread over the 12-h periods between observed perimeters in the fires analyzed. While this does not preclude the possibility that suppression can be effective over shorter timescales or under certain conditions, the results from both the full and synthetic-only models, together with the McNemar's test indicating no significant difference in predictive outcomes, suggest limited evidence within this dataset and modeling framework that suppression alone reliably alters wildfire behavior on a 12-h timeframe. These results highlight the importance of continued research into the effectiveness of retardant drops, particularly given the considerable costs of these operations and the need for a clearer understanding of their expected outcomes.

4.1. Modeling Limitations

One limitation of our modeling approach is the potential for overfitting, a known tendency of random forest classifiers. This is especially true when working with high-dimensional or imbalanced data. We took steps to reduce this risk by carefully splitting the data for training and testing the model, ensuring that all drops associated with a given fire were kept within a single split, and by focusing on balanced accuracy during evaluation. Additionally, the use of a binary classification framework introduces another source of uncertainty. The decision to classify retardant effects into two categories, along with the specific cutoff used to distinguish between them, is somewhat arbitrary and simplifies what is inherently a continuous and context-dependent process.

Another limitation of our approach is that the wildfire ROS is only calculated along the mapped fire perimeters, rather than across the entire burn area or throughout the full time between observations. As a result, fire activity occurring between perimeters—such as rapid spread during wind events or slowed progression following suppression efforts—may go undetected. This spatial and temporal discretization means our ROS estimates reflect a simplified view of fire behavior, focused on the leading edge of the fire. While this provides useful insights into perimeter-scale dynamics, it does introduce some uncertainty in interpreting ROS as a fully comprehensive measure of fire spread.

The set of available predictive features presents its own limitation. Some key drivers of fire behavior, most notably wind, were not included in our dataset due to the lack of reliable, high-resolution wind data at the spatial and temporal scales required for analysis. Wind plays a critical role in shaping fire direction, intensity, and rate of spread, and its absence likely reduced the model's ability to account for key mechanisms underlying fire growth and suppression effectiveness. Also, the environmental data incorporated into our analysis varied in spatial resolution, ranging from fine-scale (30 m) to coarse-scale (4 km).

This inconsistency may have introduced noise or bias in the feature values associated with each drop, particularly in areas with heterogeneous terrain or vegetation. As a result, some modeled relationships may reflect mismatches in data granularity rather than true ecological or suppression effects.

Some of these limitations stem from challenges inherent in quantifying fire behavior and suppression effectiveness at scale. Before settling on our final modeling framework, we explored several alternative strategies that ultimately proved unfeasible. One such approach involved predicting ROS across all points along the fire perimeters. Our goal was to build an accurate model of fire spread, which could then be used to assess the influence of retardant drops relative to expected spread patterns. However, accurately modeling ROS everywhere proved difficult, with predictions lacking sufficient accuracy. As a result, we shifted toward the counterfactual framework using synthetic drops for comparison.

We also considered a matched comparison strategy by examining areas adjacent to actual retardant drops, aiming to compare treated and untreated regions under similar conditions. While this approach offered an intuitive path to isolating suppression effects, we found it difficult to automate at scale. In many cases, multiple drops were clustered in the same area, making it difficult to find adjacent zones that were truly free of suppression. Furthermore, to ensure valid comparisons, the adjacent areas would need to be aligned with the fire front—a condition not often met given how drops are deployed in practice. These spatial and geometric complexities made it impractical to implement a systematic matching approach across the full dataset. An illustration of these challenges is shown in Figure 8.




Figure 8. The 12-h fire perimeters (shaded from white to dark red to indicate fire progression) and blue lines representing retardant drops for the 2018 Miles Fire in Oregon illustrate the spatial complexity of attempting a matched comparison strategy. Some drops are densely clustered, making it difficult to identify nearby areas that are truly untreated. Additionally, the orientation of drops is not consistently parallel to the fire front, complicating efforts to define valid comparison zones. These challenges, common across the dataset, limited our ability to implement this approach at scale.

4.2. Limitations of Interpretations

While our results suggest a lack of strong empirical support for the consistent effectiveness of aerial suppression in reducing wildfire ROS, it is important to clarify that this does not constitute evidence against its potential utility. Rather, our findings indicate that,

within the scope of our dataset and modeling framework, we were unable to isolate a clear and systematic effect of aerial retardant drops on slowing fire spread. This distinction is critical: the absence of detected impact should not be equated with proof of ineffectiveness. There are numerous potential reasons why suppression effects may not have emerged clearly in our analysis, including limitations in our ability to measure relevant fire dynamics, the granularity of available data, and the complex, context-specific goals of suppression operations.

A key consideration in interpreting suppression effectiveness is the motivation behind each flight. Aerial firefighting missions are not always intended to slow the forward spread of a fire. According to the Aerial Firefighting Use and Effectiveness (AFUE) study, suppression aircraft may be deployed to achieve a range of objectives, including reducing flame length and intensity, preventing spotting, protecting specific values at risk, or supporting ground operations [7,9]. Our dataset does not include detailed records of the operational intent behind each retardant drop, making it difficult to evaluate whether our ROS-based metric aligns with the intended outcome. Despite this limitation, approximately 60% of the drops in our dataset were carried out by large airtankers—aircraft that, according to AFUE reports, are most commonly tasked with delaying or halting the spread of the main fire front. Therefore, assessing changes in ROS remains a relevant, if imperfect, lens through which to evaluate their effectiveness.

4.3. Comparison to Existing Literature

Previous studies have provided insights into the conditions under which aerial suppression may be most effective, offering guidance on when and where these resources could be optimally deployed. For example, Wheatley et al. [11] identified environmental and operational factors that predicted initial attack success, suggesting that air resources could be more strategically allocated to fires with a higher probability of containment. Similarly, Bryan [13] found that large airtankers were particularly effective in reducing flame intensity and slowing fire growth when used in grass and shrub fuel types and in conjunction with ground efforts. These findings point to context-specific conditions that may influence suppression outcomes. In contrast, our analysis did not reveal consistent patterns linking retardant drops to reductions in fire spread. Our results do not contradict these earlier findings but rather reflect the difficulty of detecting generalized suppression effects in large, heterogeneous datasets.

Our results align more closely with those reported by Katuwal et al. [10], who found that aerial resources, specifically helicopters and air tankers, did not significantly contribute to controlled fireline length across a large sample of wildfires. While our analysis focused on ROS rather than fireline control, both suggest limited large-scale measurable impacts from aerial suppression efforts. Like Katuwal et al. [10], we also encountered the challenge of interpreting suppression effectiveness without access to detailed operational objectives. In both cases, the lack of mission-specific context limits the ability to fully assess whether the observed outcomes reflect a failure of effectiveness or a mismatch between metric and intent. The alignment between our results and those of prior studies reinforces the need for clearer operational metadata and more targeted metrics when evaluating the effectiveness of aerial suppression strategies.

5. Conclusions

This study evaluated the effectiveness of aerial retardant drops in slowing wildfire ROS using a modeling framework that incorporated both real and synthetic drop data. Despite strong model performance in predicting ROS outcomes, the results did not provide compelling evidence that aerial suppression efforts were a primary driver of reduced

fire spread rates over the 12-h periods between mapped perimeters used in this analysis. While aerial suppression may be critical, the predictive signal of whether a drop actually occurred was dominated by environmental and topographic factors in the 62 observed Oregon fires. The low importance of the real-versus-synthetic drop indicator in our full model, along with counterfactual results from the synthetic-only model, suggests that many instances of reduced ROS may have occurred independent of suppression. These findings underscore the complexity of evaluating suppression effectiveness and the limitations of using generalized metrics across large and heterogeneous datasets. Importantly, this does not imply that aerial suppression is ineffective, but rather that its measurable influence may occur over shorter timescales or in specific operational contexts that are not captured at the 12-h resolution of the available data. This highlights the need for more targeted data collection, particularly regarding operational intent, timing, and localized environmental conditions, to improve future assessments.

Building on the findings of this study, future research could benefit from key improvements in both data collection and modeling. A major need is for more detailed drop-level data, particularly information on the intended purpose of each aerial retardant drop. Understanding whether a drop was deployed to reduce fire intensity, protect assets, or slow fire spread would enable a more nuanced assessment of aerial suppression effectiveness. In addition, incorporating wind speed and direction—key drivers of fire behavior that were not available in this study—along with fire intensity data and higher-resolution versions of the existing climate datasets would enhance the model's ability to capture dynamic fire-environment interactions. A tool such as WindNinja, which generates wind fields using terrain and meteorological inputs, could be explored as a source of localized wind data for future modeling efforts. Finally, access to more frequent and accurate fire perimeter updates would allow for more precise tracking of fire spread, improving the accuracy of ROS calculations and enabling evaluation of suppression impacts at shorter temporal intervals where effects may be more evident.

Recent advances in machine learning have produced foundational models, which are systems with millions of parameters trained on terabytes of data. Presently, a set of such models has been trained using massive sets of geospatial and satellite data to classify features on the Earth's surface [23]. These models offer promising opportunities to improve future assessments of wildfire suppression effectiveness. By capturing complex spatial and temporal patterns that traditional approaches may overlook, they can support more accurate predictions of fire behavior and suppression outcomes. Leveraging transfer learning, these models can be fine-tuned for fire-specific tasks even with limited labeled data. In particular, their ability to integrate multimodal inputs—such as optical imagery, topography, and climate variables—could enhance sensitivity to localized environmental conditions, especially when high-resolution and high-frequency datasets are available. Incorporating such models into future work could provide a more nuanced understanding of the interactions between environmental factors and suppression actions across diverse fire contexts.

As wildfire activity intensifies and finite suppression resources are spread thinly over a growing fire landscape, continued efforts to improve both the quality of operational data and the sophistication of analytical tools will be essential for making informed, effective fire management decisions. This study underscores the promise of large-scale, data-driven analysis in evaluating aerial wildfire suppression tactics, while also highlighting the key challenges that must be addressed to fully harness its potential in shaping future fire management strategies.

Author Contributions: Conceptualization, L.W.-G. and J.J.; methodology, L.W.-G.; software, L.W.-G. and F.B.; validation, L.W.-G.; formal analysis, L.W.-G.; investigation, L.W.-G.; resources, F.B. and J.B.; data curation, L.W.-G., F.B. and J.B.; writing—original draft preparation, L.W.-G.; writing—review and editing, L.W.-G., J.J. and J.H.; supervision, J.J. and J.H.; project administration, J.J.; funding acquisition, J.J. and J.H. All authors have read and agreed to the published version of the manuscript.

Funding: This publication was prepared by Oregon State University using Federal funds under award #07-79-07914 from the Economic Development Administration, U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of the Economic Development Administration or the U.S. Department of Commerce.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All source code and original data used to produce the results in this study are publicly available at: <https://github.com/lindsaywiard/AerialFireFightingImpact/tree/main> (accessed on 7 May 2025).

Acknowledgments: We would like to thank Crystal Stonesifer and David Caulkin for their contribution in providing and refining the aerial retardant drop dataset used in this study. Their work to curate and link aerial drop records to wildfire incidents made this analysis possible.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AFUE	Aerial firefighting use and effectiveness
ATU	Additional telemetry unit
AUS	Aviation use summary
FEDS	Fire events data suite
GIS	geographic information system
IA	Initial attack
MTBS	Monitoring trends in burn severity
ROS	Rate of spread
USFS	United States Forest Service
VIIRS	Visible Infrared Imaging Radiometer Suite

References

1. Iglesias, V.; Balch, J.K.; Travis, W.R. U.S. fires became larger, more frequent, and more widespread in the 2000s. *Sci. Adv.* **2022**, *8*, eabc0020. [\[CrossRef\]](#) [\[PubMed\]](#)
2. McGinnis, S.; Kessenich, L.; Mearns, L.; Cullen, A.; Podschwit, H.; Bukovsky, M. Future regional increases in simultaneous large Western USA wildfires. *Int. J. Wildland Fire* **2023**, *32*, 1304–1314. [\[CrossRef\]](#)
3. Calkin, D.E.; Stonesifer, C.S.; Thompson, M.P.; McHugh, C.W. Large airtanker use and outcomes in suppressing wildland fires in the United States. *Int. J. Wildland Fire* **2014**, *23*, 259–271. [\[CrossRef\]](#)
4. Stonesifer, C.S.; Calkin, D.E.; Thompson, M.P.; Belval, E.J. Is This Flight Necessary? The Aviation Use Summary (AUS): A Framework for Strategic, Risk-Informed Aviation Decision Support. *Forests* **2021**, *12*, 1078. [\[CrossRef\]](#)
5. Heines, B.; Lenhart, S.; Sims, C. Assessing the economic trade-offs between prevention and suppression of forest fires. *Nat. Resour. Model.* **2018**, *31*, e12159. [\[CrossRef\]](#)
6. Schammel, M.H.; Gold, S.J.; McCurry, D.L. Metals in Wildfire Suppressants. *Environ. Sci. Technol. Lett.* **2024**, *11*, 1247–1253. [\[CrossRef\]](#) [\[PubMed\]](#)
7. Plucinski, M.P. Fighting Flames and Forging Firelines: Wildfire Suppression Effectiveness at the Fire Edge. *Curr. For. Rep.* **2019**, *5*, 1–19. [\[CrossRef\]](#)

8. Plucinski, M.P.; Pastor, E. Criteria and methodology for evaluating aerial wildfire suppression. *Int. J. Wildland Fire* **2013**, *22*, 1144–1154. [[CrossRef](#)]
9. United States Forest Service. *Aerial Firefighting Use and Effectiveness (AFUE) Report*; Technical Report; U.S. Department of Agriculture, Forest Service: Washington, DC, USA, 2020.
10. Katuwal, H.; Calkin, D.E.; Hand, M.S. Production and efficiency of large wildland fire suppression effort: A stochastic frontier analysis. *J. Environ. Manag.* **2016**, *166*, 227–236. [[CrossRef](#)]
11. Wheatley, M.; Wotton, B.M.; Woolford, D.G.; Martell, D.L.; Johnston, J.M. Modelling initial attack success on forest fires suppressed by air attack in the province of Ontario, Canada. *Int. J. Wildland Fire* **2022**, *31*, 774–785. [[CrossRef](#)]
12. Stonesifer, C.S.; Calkin, D.E.; Thompson, M.P.; Stockmann, K.D. Fighting fire in the heat of the day: An analysis of operational and environmental conditions of use for large airtankers in United States fire suppression. *Int. J. Wildland Fire* **2016**, *25*, 520–533. [[CrossRef](#)]
13. Bryan, C.R. Evaluating the Efficiency, Equity, and Effectiveness of Wildfire Suppression Strategy Using the Microeconomic Toolkit. Ph.D. Thesis, Colorado State University, Fort Collins, CO, USA, 2024.
14. Monitoring Trends in Burn Severity. Available online: <https://www.mtbs.gov/> (accessed on 13 March 2025).
15. U.S. Geological Survey. *Elevation Derivatives for National Applications*; Fact Sheet 2005-3049; U.S. Geological Survey: Washington, DC, USA, 2005. [[CrossRef](#)]
16. Abatzoglou, J.T. Development of gridded surface meteorological data for ecological applications and modelling. *Int. J. Climatol.* **2013**, *33*, 121–131. [[CrossRef](#)]
17. Wang, T.; Hamann, A.; Spittlehouse, D.; Carroll, C. Locally Downscaled and Spatially Customizable Climate Data for Historical and Future Periods for North America. *PLoS ONE* **2016**, *11*, e0156720. [[CrossRef](#)] [[PubMed](#)]
18. Thornton, M.M.; Shrestha, R.; Wei, Y.; Thornton, P.E.; Kao, S.C. *Daymet: Monthly Climate Summaries on a 1-km Grid for North America, Version 4*; Technical Report; ORNL DAAC: Oak Ridge, TN, USA, 2020. [[CrossRef](#)]
19. Rollins, M.G. LANDFIRE: A nationally consistent vegetation, wildland fire, and fuel assessment. *Int. J. Wildland Fire* **2009**, *18*, 235–249. [[CrossRef](#)]
20. Jones, M.O.; Robinson, N.P.; Naugle, D.E.; Maestas, J.D.; Reeves, M.C.; Lankston, R.W.; Allred, B.W. Annual and 16-Day Rangeland Production Estimates for the Western United States. *Rangel. Ecol. Manag.* **2021**, *77*, 112–117. [[CrossRef](#)]
21. Vermote, E. *NOAA Climate Data Record (CDR) of AVHRR Normalized Difference Vegetation Index (NDVI), Version 5*; NOAA National Centers for Environmental Information: Boulder, CO, USA, 2019. [[CrossRef](#)]
22. Maskey, M. NASA Earthdata VEDA Dashboard Fire Perimeters. 2025. Available online: <https://www.earthdata.nasa.gov/dashboard/data-catalog/fire> (accessed on 30 January 2025).
23. Leese, L. Expanded AI Model with Global Data Enhances Earth Science Applications. NASA, 4 December 2024. Available online: <https://science.nasa.gov/science-research/ai-geospatial-model-earth/> (accessed on 30 April 2025).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.