Fighting Wildfires under Uncertainty: A Sequential Resource Allocation Approach*

Hau Chan[1] , **Long Tran-Thanh**[2] and **Vignesh Viswanathan**[3]
[1]University of Nebraska-Lincoln, USA
[2]University of Southampton, UK
[3]Indian Institute of Technology, Kharagpur, India
hchan3@unl.edu, l.tran-thanh@soton.ac.uk, vigneshviswanathan@iitkgp.ac.in

Abstract

Standard disaster response involves using drones (or helicopters) for reconnaissance and using people on the ground to mitigate the damage. In this paper, we look at the problem of wildfires and propose an efficient resource allocation strategy to cope with both dynamically changing environment and uncertainty. We propose Firefly, a new resource allocation algorithm, that can provably achieve optimal or near-optimal solutions with high probability by first efficiently allocating observation drones to collect information to reduce uncertainty, and then allocate the firefighting units to extinguish the fire. For the former, Firefly uses a combination of maximum set coverage formulation and a novel utility estimation technique, and it uses a knapsack formulation to calculate the allocation for the latter. We also demonstrate empirically by using a real-world dataset that Firefly achieves up to $80-90\%$ performance of the offline optimal solution, even with a small number of drones, in most cases.

1 Introduction

In recent years, forest fires pose an increasing challenge for the local and statewide government agencies around the globe due to human and natural causes (e.g., effects of climate change). Over the last ten years in the United States, it is estimated that 54% and 46% of the wildfires are caused by human and lightning.With growing forest fire sizes and instances and demand for suppression resources, fighting forest fires has developed into a major budgetary concern for many nations. In the United States alone, the National Interagency Fire Center (NIFC) reports that in 2017, 71,499 fires burned a total of 10,026,086 forested acres resulting in a suppression cost of $2,918,165,000, making 2017 the most expensive year on record[1]. In 2018, according to the California Department of Forestry and Fire Protection (Cal Fire) and the NIFC, the infamous California wildfire season had 8,527 fires with a total burning area of 1,893,913 acres resulting in an operational cost of $432 million and insurance claims of $12 billion. Historical (NIFC) data on the costs of forest fires reveals that both the number of acres being burned and the number of suppression resources utilized to suppress fires are rising. *How can we utilize these resources to suppress fires in the most effective way? How can we dispatch the resources appropriately to different zones?* These are the questions we aim to address in this paper.

According to the United States Forest Service, the resource dispatch processes begins when a fire center receives reported incidents or sighting. The incidents are reported to the corresponding center that manages the surrounding areas or dispatch areas. The center divides its dispatch areas into smaller dispatch zones and determines the types of resources to be dispatched depending on various conditions. The fire conditions are categorized into low, moderate or high dispatch levels. The resources dispatched to a particular zone depend on the zone's value, which depends on various factors and fire conditions. In addition, each zone's value encompasses some measures of life, private property, critical infrastructure, cultural/natural resources, and other factors. When there are multiple simultaneously fires in the dispatch areas, the fire center then needs to prioritize the zones and resources. The resources are usually shared between all other (local, state, or federal) agencies that are responsible for the areas so that the fire can be managed more efficiently.

Various resources have been used to aid in the fight against forest fires. In recent years, agencies such as the Los Angeles Fire Department, the United States Department of Agriculture, and the United States Forest Service have been combining drones with more traditional units to help combat wildfire[2]. The primary purpose of drones is to provide the agencies or firefighters information (i.e., the fires and the surrounding areas). As a result, we can characterize the resources into two main types: surveillance and fire fighting resources. For simplicity, we will call the surveillance resources drones and the firefight resources firefighters. As dis-

*The authors equally contributed to the paper.

[1]National Interagency Fire Center. Historical Wildland Fire information: Suppression Costs". https://www.nifc.gov/fireInfo/fireInfo_documents/SuppCosts.pdf

[2]https://www.latimes.com/local/lanow/la-me-ln-lafd-drone-skirball-fire-20171214-story.html
https://www.nbcnews.com/mach/science/drones-are-fighting-wildfires-some-very-surprising-ways-ncna820966
https://www.usda.gov/media/blog/2019/07/09/drones-provide-eye-sky-help-fight-fires
https://www.fs.fed.us/managing-land/fire/aviation/uas/

cussed, drones provide us information about the zones and fires such as the strength of the fire in the deployed zones, and firefighters combat the fire in the deployed zones. Note that the latter can be either humans or machines.

It is not hard to see that if the number of firefighters is sufficient for a zone, we can extinguish the fire. However, both drones and firefighters are limited, and we must make some trade-off and prioritize when we don't have enough resources for fires in all of the zones with fires. In addition, when a fire breaks out in some zones of the forest, we do not a *prior* know the strength of the fires. As a result, determining the number of resources (e.g., firefighters and drones) needed to put out a fire in any zone is a necessary and challenging problem. Note that the strength of a wildfire in a zone changes dynamically over time (e.g., days, months, seasons, and years). Even if we know its strength in a particular time step, we are sure to be uncertain in the future time step. Finally, we need to be able to dispatch resources to fight fires in multiple zones, which is a challenging problem in itself. As a result, to successfully fight a wildfire, we need to address the above challenges and *learn* about the strength of the fires over time and *allocate* resources to fight the fires.

To tackle the domain challenges in fighting a forest fire, we propose Firefly, a novel algorithm to allocate a limited quantity of both drones and firefighters among several zones. In particular, this algorithm first formulates the drone allocation as a maximal set coverage problem, and uses a greedy algorithm to allocate, and therefore, observe both the fire strength and value of a subset of zones. It then calculates an unbiased estimator of the fire strength and value of the remaining zones. Using both the observed and estimated data, Firefly then solves the firefighter allocation problem as a knapsack. We prove that if the number of drones is $\Theta(\ln^2(n))$ where n is the number of zones, and the topology of the zones (i.e., how the zones are adjacent to each other - see Section 2 for the precise definition) follows a random graph structure, then Firefly can solve the firefighter allocation problem optimally with high probability. Furthermore, we show that if the topology of the zones is an arbitrary graph, but the total amount of environmental changes within T time steps (see Section 3.5 for a more precise definition of environmental changes) is limited by a budget B, then a modified version of Firefly can achieve an additive approximation of $\tilde{O}(n^{1/3}B^{1/3}T^{2/3})$, compared to the offline optimal solution, where $\tilde{O}$ hides away all the logarithmic dependencies. Finally, we numerically evaluate the performance of Firefly in a large variety of settings. Our results show that Firefly, when equipped with a relatively small number of drones (e.g., $10-20\%$ of the total number of zones), can achieve up to $80-95\%$ of the optimal performance in most cases.

1.1 Related Work

In this section, we briefly cover related work on resource allocation for wildfire mitigation. From the perspective of wildfire modeling, [Bendix and Commons, 2017] study the distribution and frequency of wildfire in California riparian ecosystems. [Malamud *et al.*, 2005] present a methodology for characterising fire regimes based on past data. [Salis *et al.*, 2012] study the behaviour of wildfire in North Sardinia, Italy. [Piñol *et al.*, 1998] and [Westerling *et al.*, 2006] study the effect of climate change on wildfires in different regions. While these works help us learn more about the nature of wildfire, none of them address the problem of resource allocation or uncertainty. To capture the uncertainty of the nature of wildfire, [Cruz *et al.*, 2008] develop a model to predict wildfire behaviour in pine plantations. [Rodriguez Aseretto *et al.*, 2013] and [de Rigo *et al.*, 2013] propose data-driven models to predict wildfire behaviour. [Yanjun Li *et al.*, 2006] and [Sabit *et al.*, 2011] propose wireless sensor network based wildfire detection and monitoring systems. [Boulton *et al.*, 2016] propose using social media to detect and map wildfires. [Kyzirakos *et al.*, 2014] use satellite images to monitor wildfire. While these works do reduce the uncertainty involved in wildfires, they do not propose any algorithm to allocate resources efficiently. From the resource allocation perspective, [Donovan and Rideout, 2003; Hu and Ntaimo, 2009; Ntaimo *et al.*, 2012; Lee *et al.*, 2012] all propose algorithms to deploy firefighters to suppress wildfires. Our work differs from theirs in two aspects. First, all of their work assumes knowledge about the behaviour of the wildfire. Second, they ignore the exploration aspect of this problem and only consider the allocation of firefighting resources. [Belval *et al.*, 2014] and [HomChaudhuri and Cohen, 2010] propose algorithms to efficiently place firefighters within a region so as to maximize the fire which gets suppressed. This is a parallel problem and their work can be used along with ours. [Tsang *et al.*, 2013] study resource sharing among different fire departments. However, they do not provide an algorithm to allocate these resources efficiently. Finally, a number of works have focused on efficient drone allocation (e.g., [Chmaj and Selvaraj, 2015; Yuan *et al.*, 2015]). However, these works do not take into account the problem of firefighting resource allocation and thus cannot be applied to our setting.

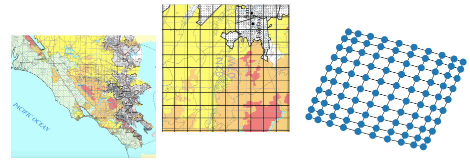

(a) (b) (c)

Figure 1: (a) Fire Hazard Severity Zones of Marin County, CA (Yellow = Moderate, Orange = High, Red = Very High), (b) 10 x 10 Zoning Example (of a sub-zone), (c) Undirected Grid Graph Representation of the 10 x 10 Zones.

2 Problem Description

We formulate our optimization problem as follows: Consider a set of forest zones $N = \{1, 2, \ldots, n\}$, which are vulnerable to catching fire. While there exist various definitions of vulnerable, the most known or common vulnerability measure is the probability of ignition. The probability of ignition is the chance that a firebrand will cause an ignition when it is combined with fuels. The vulnerability or the probability of ignition of a zone typically depends on the current temperature, shading from forest canopy or cloud cover, and 1-hr fuel moisture content [Andrews, 2014] (see also National Wildfire Coordinating Group.) We allocate to each of these zones a priority value $k_i, \forall i \in N$, that captures the importance of the zone. That is, how valuable is it for us to protect that particular zone? The importance of a zone or value can be defined appropriately based on the particular ecosystems (e.g., trees, species, and lands) in the zone.

We use the directed graph G_R to represent the topology of these zones: Each zone is a node and a directed edge connects node i with node j if node j can be observed from node i. That is, if (i, j) is an edge of G_R, then, by deploying a drone at i, we will obtain information (e.g., the fire strength) about both i and j. Note that this graph is useful for drone deployment (see Section 3.1 for more details). Figure 1 presents an illustration of the representation of Marin County, CA [3].

For each time step $t = \{1, 2, \ldots, T\}$, zone i can catch fire with strength $0 \leq a_{i,t} \leq C$ (with $a_{i,t} = 0$ indicating there is no fire). In Ecology, the strength of fire can be associated with fire intensity [Keeley, 2009], which can be described and measured as fire energy released, for instance, using fireline intensity [Keeley, 2009]. Roughly speaking, fireline intensity is the rate of heat transfer per unit length of the fire line. Therefore, fireline intensity is a good measure of the difficulties of suppressing fire and the number of firefighting resources required to put out the fire with a certain intensity. As a result, we associate the strengths (i.e., $a_{i,t}$) as the number of firefighter units required to extinguish the fire. It is also reasonable to assume that such upper bound value C exists. A possible justification for this is, e.g., the fact that the strength of fire (e.g., fire intensity) could be limited by the total number of trees within the zone, which is finite. We consider a realistic setting where $a_{i,t}$ is unknown to us *a priori*, and can only be revealed if at least one firefighter unit is allocated to zone i, or the zone is covered by at least one observer drone at time step t (see later for more details). This is typically the setting where some fire has started in some zones, and the fire incidents are reported to agencies, which will determine the resources to deploy to zones depending on the fire's strengths.

In our setting, we have r firefighters (who can be both machines such as drones or ground vehicles, and human units) and m observer drones. To extinguish the fire at zone i at time step t, we need to allocate sufficient numbers of firefighters to zone i. We let $d_{i,t}$ to be the firefighter units allocated to zone i at time t. The reward of this allocation or action is given by

$$f_i(d_{i,t}) = \begin{cases} k_{i,t} & d_{i,t} \geq a_{i,t} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

That is, we only receive reward $k_{i,t}$ if there is a fire in zone i and we have deployed a sufficient amount of firefighters in that zone. The value of $k_{i,t}$ can vary over time. For example, partially burnt zones might lose their values in the future, or due to some seasonal changes, some zones can become more important. We obtain the following optimization problem:

$$\begin{aligned} &\max_{\{d_{i,t}\}} \quad \sum_{t=1}^{T} \sum_{i=1}^{n} f_i(d_{i,t}) \\ &\text{subject to} \quad \forall t: \ \sum_{i=1}^{n} d_{i,t} \leq r, \quad \forall i, t: \ d_{i,t} \in \mathbb{N} \end{aligned} \tag{2}$$

That is, our goal is to identify the optimal allocation of firefighter units that maximises the total rewards[4]. If we know the exact value of $a_{i,t}$ fire strength and $k_{i,t}$, the problem can be reduced to a $0-1$ knapsack problem with bounded weights (as fire strengths are bounded above by C), which can be optimally solved within running time complexity polynomial in the number of zones [Pisinger, 1999]. However, as described earlier, the exact value of $a_{i,t}$ and $k_{i,t}$ is not known unless a drone covers zone i or a firefighter unit is sent to that zone. Thus, our model consists of the following practical settings:

- At each time step t, each zone i has a pair of values $k_{i,t}$ and $a_{i,t} \geq 0$, which are not fully observed unless we allocate firefighters and/or drones to cover them.
- If we allocate a drone to a specific zone, we assume it gives us information regarding the strength of fire at that particular zone and at all the zones adjacent to it in G_R.
- If we allocate a certain number of firefighters ($d_{i,t} > 0$) to zone i at time step t, even if we do not receive any positive utility (i.e., $d_{i,t} < a_{i,t}$), we assume that we get to know the strength of fire ($a_{i,t}$) in that zone.

Regarding the nature of uncertainty about $a_{i,t}$ and $k_{i,t}$, we consider the worst-case scenario, that is, the sequences of $a_{i,1}, \ldots a_{i,T}$ and $k_{i,1}, \ldots k_{i,T}$ *can arbitrarily change over time*. Note that in practice, both $a_{i,t}$ and $k_{i,t}$ can be dictated by some (hidden) processes. For example, if the time step is relatively rapid (e.g., it represents days), then if a fire is not extinguished yet, its strength will be increased over time (up to upper bound C). On the other hand, if time step represents seasons, then it could be the case that $a_{i,t}$ does not change much over time (e.g., as fire might be caused and governed by nature and its laws). However, to make our model generic (and thus, to be able to cover as many different scenarios as possible), we do not impose any further assumptions on the nature of $a_{i,t}$ and $k_{i,t}$. To reduce the uncertainty of the fire strengths, we can deploy drones to observe the zones. Let $v_t \in \{0, 1\}^n$ denote the indicator vector where $v_t(i) = 1$ if a drone is allocated to zone i at time step t. In addition, let $\mathcal{N}(v_t) \in \{0, 1\}^n$ denote the neighbourhood of v_t in G_R. That is, $\mathcal{N}(v_t)(i) = 1$ if and only if $v_t(i) = 1$ or there exists $j \neq i$ such that (i, j) is an edge in G_R and $v_t(j) = 1$. Note that in our setting, any zone within $\mathcal{N}(v_t)$ is observable (i.e., we can observe the fire's strength in that zone).

Let $\mathcal{D}(v_t)$ denote the set of firefighter allocation policies that take into account the information received from drone

[3] https://osfm.fire.ca.gov/divisions/wildfire-prevention-planning-engineering/wildland-hazards-building-codes/fire-hazard-severity-zones-maps/

[4] We can also formulate our model as a loss/damage minimization problem. In particular, $f_i(d_{i,t})$ can be replaced with a loss function $l_i(d_{i,t})$, which provides a positive loss k_i if $a_{i,t} > 0$ and $d_{i,t} < a_{i,t}$. The reason we formulate our model as a maximization problem is that the underlying optimization problem is the same, namely the knapsack model. Using the loss minimization model would need an extra conversion step to fit it into the knapsack setting.

allocation v_t. The objective in Eq (2) can be rewritten as:

$$\max_{\{v_t\},\{d_{i,t}\}\in\mathcal{D}(v_t)} \quad \sum_{t=1}^{T}\sum_{i=1}^{n} f_i(d_{i,t})$$
$$\text{subject to} \qquad \forall t: \ \sum_{i=1}^{n} d_{i,t} \le r, \ \ \forall i,t: \ \ d_{i,t}\in\mathbb{N} \tag{3}$$

Our goal thus is to identify efficient drone deployments and their corresponding firefighter allocations that maximizes the total utility value over T time steps.

3 The Firefly Algorithm

Given the problem description above, we now turn to the design of our algorithm called Firefly. Note that while both the drone and firefighter allocation tasks can be combined into a single optimization problem, we follow the route of separating these problems from each other. The reason for this decoupling is the fact that we can significantly reduce the search space (both problems are difficult combinatorial problems with large search spaces). As such, the high-level workflow of Firefly can be described as follows. At each step t:

1. Find a drone allocation as a solution of a maximum set coverage problem (see Section 3.1);
2. Use the learnt information from the drone allocation to estimate the utility function of each zone (Section 3.2);
3. Use the abovementioned estimates to solve the firefighter allocation as a knapsack problem (Section 3.4).

3.1 Drone Allocation as Max Coverage

The goal of this phase is to reduce the uncertainty about as many values of $a_{i,t}$ (i.e., the strength of fire at zone i at time step t) as possible. More importantly, we are interested in covering (i.e., identifying the correct fire strength) of the zones where there is fire (i.e., $a_{i,t} > 0$), while avoiding sending our drones to zones without fire (to not waste our resources). As such, our goal is to solve the following optimization problem for each time step t (using the notations from Section 3): $\max_{\{v_t\}} \sum_{i=1}^{n}\Big(\mathbb{I}\{a_{i,t} > 0, \ N(v_t)(i) = 1\}\Big)$ where $\mathbb{I}\{\}$ is the indicator function. That is, we want to identify the v_t vector of drone allocation that maximizes the number of zones on fire and are within the observation range (i.e., neighbourhood) of at least one drone. We consider two cases: (i) we know the zones with fire in advance (but not their strength); and (ii) we do not have any information about the locations of fire. The first case deals with the setting when we receive a detail report of fire while the second case deals with general fire or smoke sightings and reporting. Ideally, we would like to cover the whole graph. However, due to the limited number of drones, we are not able to cover the whole graph. Below, we discuss how we can address this challenge for the above cases.

We tackle the former first. In particular, let $u_t \in \{0,1\}^n$ denote the position vector of fire (i.e., $u_{i,t} = 1$ if and only if $a_{i,t} > 0$). We construct $\hat{G}_t$ graph from G_R and u_t as follows:

1. For each node i of G_R we create a copy i^1 in $\hat{G}_t$;
2. For each j with $u_{j,t} = 1$ create additional copy j^2 in $\hat{G}_t$;
3. For each pair of i^1 and j^2, add edge (i^1, j^2) to $\hat{G}_t$ if there is edge going from i to j in G_R.

We can show that the optimization problem described above is equivalent to the maximum set coverage problem. More specifically, in the maximum set coverage problem, we are given a collection of sets $S = \{S_1, ..., S_m\}$ and an integer k, we want to find $T \subseteq \{1, ..., m\}$ of size k such that $|\cup_{i\in T} S_i|$ is maximized. To reduce our problem to the maximum set coverage problem, for each node i^1 of $\hat{G}_t$, we construct a set of nodes that i^1 is connected to in $\hat{G}_t$. For each node j^2 of $\hat{G}_t$, we construct a set containing only j^2. Since we have m drones, we want to find m sets (or select m nodes) as to cover as many of the fire locations as possible so that we can determine their strengths. While the maximum coverage problem is known to be NP-hard, we can apply an efficient greedy algorithm to find a solution that obtains at least $(1 - 1/e)$ fraction of the optimal solution [Hochbaum and Pathria, 1998]. The greedy algorithm iteratively (until we have selected k sets) adds a set that provides the maximal marginal increase in the number of elements covered given the already selected sets. We incorporate such greedy algorithm to find v_t in our Firefly algorithm. When we do not have access to vector u_t, we simply solve the maximum set coverage problem on G_R where the set of each node i in G_R is defined to be the neighbors it covers including i. For the sake of simplicity, we say that $\hat{G}_t = G_R$. Firefly then uses the same greedy algorithm to allocate the drones on $\hat{G}_t$.

Algorithm 1: The Firefly Algorithm

Inputs: G_R, r, m, C, η
for *t = 1, . . . , T* **do**
  Environment generates fire strength vector a_t and reward vector k_t
  $u_t \leftarrow$ Fire location vector (if available)
  $v_t \leftarrow$ Greedy max-m coverage on $\hat{G}_t$ (Sec 3.1)
  Estimate $\hat{a}_{i,t}$ and $\hat{k}_{i,t}$ (Sec 3.2)
  $w^*_{i,t} \leftarrow$ Solve perturbed knapsack from Eq (7)
  $d_{i,t} = \lceil \hat{a}_{i,t} \rceil w^*_{i,t}$
  $\forall i$: Allocate $d_{i,t}$ firefighters to i and observe true reward $f_t(d_{i,t})$ (Sec 3.4)
  Update $\hat{A}_{i,t}$ and $\hat{K}_{i,t}$ (Sec 3.2 and 3.3)
end

3.2 Utility Value Estimation

Having the drones allocated, we need to gather the information about the true values of observed $a_{i,t}$ and $k_{i,t}$, and provide estimates for the unobserved ones, as this is necessary to calculate the utility function of each zone i. This part of the algorithm is particularly important because not every zone is observed from the drone allocation stage. We start with the estimation of the fire strength. To do so, we maintain two parameters for each zone, $\hat{a}_{i,t}$ for the estimated value of the current fire strength (i.e., $a_{i,t}$), and $\tilde{A}_{i,t}$ is the cumulative fire strength at zone i up to time step t. Similarly, we maintain two parameters, $\hat{k}_{i,t}$ for the estimated value of the current reward $k_{i,t}$, and $\tilde{K}_{i,t}$ is the cumulative reward at zone i up to time step t. We estimate $\hat{a}_{i,t}$ as follows: If i is covered, $\hat{a}_{i,t} = a_{i,t}$ (as the true value of fire strength is observed). Otherwise, $\hat{a}_{i,t} = \frac{1}{t-1}\tilde{A}_{i,t-1}$ for $t > 1$, and $\hat{a}_{i,1} = 0$. While $\hat{k}_{i,t}$ is estimated as: If i is covered, $\hat{k}_{i,t} = k_{i,t}$ (as the true reward is observed). Otherwise, $\hat{k}_{i,t} = \frac{1}{t-1}\tilde{K}_{i,t-1}$ for $t > 1$, and $\hat{k}_{i,1} = 0$. The estimated utility function for each zone will be $\hat{f}_{i,t}(d) = \hat{k}_i$ if $d \geq \hat{a}_{i,t} > 0$, and 0 otherwise. We then use these utility estimates to solve the firefighter allocation problem (see Section 3.4 for more details). After the allocation is done, and information about the true utility values and fire strength have been collected, we finish our estimation phase by updating $\tilde{A}_{i,t}$ and $\tilde{K}_{i,t}$ as follows: If i is covered: $\tilde{A}_{i,t} = \tilde{A}_{i,t-1} + a_{i,t}$ and $\tilde{K}_{i,t} = \tilde{K}_{i,t-1} + k_{i,t}$. Otherwise let $\hat{p}_{i,t}(d)$ denote the estimate of the probability that the firefighter allocation algorithm (which can provide a randomized allocation), described later in Section 3.4, allocates exactly d resources to zone i at time step t. We have

$$\tilde{A}_{i,t} = \tilde{A}_{i,t-1} + a_{i,t} \sum_{d>0} \frac{\mathbb{I}\{d_{i,t} = d\}}{\hat{p}_{i,t}(d)} \quad (4)$$

$$\tilde{K}_{i,t} = \tilde{K}_{i,t-1} + \sum_{d>0} f_{i,t}(d) \frac{\mathbb{I}\{d_{i,t} = d\}}{\hat{p}_{i,t}(d)} \quad (5)$$

The intuition behind Eqs (4) and (5) is that if we know $p_{i,t}(d)$, the true probability that $d_{i,t} = d$, then replacing $p_{i,t}(d)$ into the second term on the RHS of Eqs (4) and (5) gives an unbiased estimate of $a_{i,t}$ and $k_{i,t}$, whenever $d_{i,t} = d$. As such, we can use it to estimate the strength of fire and reward at zone i in many cases without being able to observe their true value, as in expectation the unbiased estimates will return the true values. Note that this technique is common in the online adversarial learning literature for estimating unseen loss (for more details see, e.g., [Neu and Bartók, 2013]). Now, the main challenge is that the true probability $p(d_{i,t} = d)$ cannot be explicitly calculated in many situations (e.g, then the underlying firefighter allocation algorithm needs to solve a complex optimization problem, as it is the case in our setting). To overcome this issue, we discuss a geometric resampling based technique to calculate $\hat{p}_{i,t}(d)$, an estimate of $p(d_{i,t} = d)$.

3.3 Calculating $\hat{p}_{i,t}(d)$ with Geometric Resampling

In this section, we provide a technique to estimate $p(d_{i,t} = d)$, the probability that the optimizer F, described later in Section 3.4, would have allocated d resources to zone i. This probability value is typically unknown, as we can only observe a concrete outcome of F (i.e., a particular allocation $\{d_{i,t}\}$). This probability can be estimated by using a technique called geometric resampling [Neu and Bartók, 2013], which lies on the following simple idea: if we repeatedly use F to reallocate the resources for many rounds, and observe whether the allocation at zone i is d at each round, then this observation follows a geometric distribution. Therefore we can use geometric distribution to infer $p(d_{i,t} = d)$.

Algorithm 2: The Geommetric Resampling Algorithm

Inputs: drone allocation v; $L_{\max} > 0$, optimizer F
$\forall i$ zone: $L_i(d) = L_{\max}$
for $l = 1, 2, \ldots, L_{\max}$ **do**
  Recalculate firefighter allocation $\hat{d}_{i,t}$ using F and v
  for *all zone* i **do**
    If $(N(v)(i) \neq 1$ and $\hat{d}_{i,t} = d)$ then $L_i(d) =$ `min(`$L_i(d)$, l`)`
  end
end
return $\{L_i(d)\}_{i,d}$

More specifically, the adaptation of this technique to our setting is given in Algorithm 2. For the sake of completeness we explain the algorithm in detail as follows:

- Step 1: Use F to repeatedly calculate the firefighter allocation for sufficient number of times (Step 3);
- Step 2: Let $L_i(d)$ denote the first recalculation round in which $d_{i,t} = d$ (i.e., F allocates d resources to zone i). It is easy to prove that $L_i(d)$ follows a geometric distribution with $\mathbb{E}[L_i(d)] = \frac{1}{p(d_{i,t}=d)}$. As such, $1/L_i(d)$ is a unbiased estimator of $p(d_{i,t} > d)$.
- Step 3: However, $L_i(d)$ can be undefined as it can grow to infinity. Therefore, by only recalculating up to $L_{\max}$ rounds (i.e., $L_i(d)$ is capped by $L_{\max}$) and setting $\hat{p}_{i,t}(d) = 1/L_i(d)$, we can still get a good (but biased) estimate of $p(d_{i,t} = d)$.

3.4 Firefighter Allocation as Knapsack Problem

Given the description of drone allocation and utility estimation, we now turn to the discussion of the firefighter allocation phase. In particular, taking into account the fire strength vector $\hat{a}_t$, we solve the following knapsack problem:

$$\begin{aligned} \max_{\{w_{i,t}\}} \quad & \sum_{i=1}^{n} \hat{k}_{i,t} w_{i,t} \\ \text{subject to} \quad & \sum_{i=1}^{n} \lceil \hat{a}_{i,t} \rceil w_{i,t} \leq r, \;\; \forall i: \;\; w_{i,t} \in \{0,1\} \end{aligned} \quad (6)$$

where $w_{i,t}$ specifies whether $\lceil \hat{a}_{i,t} \rceil$ firefighters are deployed to zone i at time t. The intuition behind this formulation is that from a utility perspective, it is not useful to allocate less than $\hat{a}_{i,t}$ resources to zone i, as we cannot receive any rewards for doing so. Therefore, if we decide to allocate resources to that zone, it is more beneficial to allocate at least $\lceil \hat{a}_{i,t} \rceil$ (as $\hat{a}_{i,t}$ can be fractional). However, note that $\hat{a}_{i,t}$ is not always the true fire strength and $\hat{k}_{i,t}$ is not always the true reward. As such, to capture the uncertainty around the true value of $\hat{a}_{i,t}$ and $\hat{k}_{i,t}$, we perturb the knapsack problem in Eq (6) by adding some extra random noise to the model to allow a certain degree of exploration, as opposed to solely focusing on exploitation (i.e., assuming $\hat{a}_{i,t}$ and $\hat{k}_{i,t}$ are true values and solve the knapsack). Another reason to perturb the knapsack problem is to increase the number of zones with positive $P(d_{i,t} > 0)$. Recall that if $P(d_{i,t} > 0) > 0$ then the

RHS of Eqs (4) and (5) give us a good estimator for the true value of $a_{i,t}$ and $k_{i,t}$ (see Section 3.2 for more details).

Given this, for each i, let $z_{i,t} \sim Exp(\eta)$ denote a random variable sampled from an exponential distribution with parameter η. Now, we define $\hat{k}_{i,t}$ the perturbed value of k_i as follows: If $N(v_t)(i) = 1$ (i.e., i is covered by a drone), $\tilde{k}_{i,t} = k_{i,t}$. Otherwise $\tilde{k}_{i,t} = \hat{k}_{i,t} + z_{i,t}$. We then solve the following perturbed knapsack problem at each time step t:

$$\begin{aligned} \max_{\{w_{i,t}\}} \quad & \sum_{i=1}^{n} \tilde{k}_{i,t} w_{i,t} \\ \text{subject to} \quad & \sum_{i=1}^{n} \lceil \hat{a}_{i,t} \rceil w_{i,t} \leq r, \ \ \forall i: \ \ w_{i,t} \in \{0,1\} \end{aligned} \tag{7}$$

The firefighter allocation then can be calculated as follows: We allocate $\lceil \hat{a}_{i,t} \rceil w^*_{i,t}$ to each zone i, where $w^*_{i,t}$ is a solution of the knapsack problem in Eq 7. Note that as $0 \leq a_{i,t} \leq C$, $\lceil \hat{a}_{i,t} \rceil$ are also bounded above by C. Therefore, if C is independent from r and other parameters, this knapsack problem can also be solved in polynomial time (by using, e.g., dynamic programming [Pisinger, 1999]). Given this allocation, the total reward at time t will be $R_t = \sum_{i=1}^{n} f_i(\lceil \hat{a}_{i,t} \rceil w^*_{i,t})$.

3.5 Performance Analysis of Firefly Algorithm

Based on the description of each phase in Sections 3.1 - 3.4, the complete Firefly algorithm is depicted in Algorithm 1. As we do not have any restrictions on how $a_{i,t}$ and $k_{i,t}$ can change over time, it is difficult in general to provide efficient performance guarantees for the algorithm. In particular, it is easy to show that this problem is APX-hard, by e.g., considering the case of $m = 0$ (i.e., no drones for observation), and $a_{i,t}$ and $k_{i,t}$ change arbitrarily over time (i.e., no algorithm can do better than a uniform random allocation). However, under some additional conditions, we can achieve theoretical guarantees. In particular, our first result shows that if we have $O(\ln^2(n))$ number of drones, and G_R follows certain randomized structure, then Firefly can achieve optimal solution (i.e., maximal total reward) with high probability. More precisely:

Theorem 1. *Suppose that G_R consists of a finite number of Erdos-Renyi random graphs (we also allow extra edges between these graphs). Let $p > 0$ be the smallest edge-generation probability of these graphs. Then for each $p > 0$ and $\delta > 0$, if $m \in \Theta(\ln^2(n))$, Firefly finds the optimal solution with at least $(1-\delta)$ probability.*

4 Experimental Results

Our goal is to evaluate the performance of our Firefly algorithm under different parameter settings and zoning/graph structures of certain areas. Recall that the Firefly algorithm takes G_R (e.g., the zoning graph), r (e.g., the number of firefighters), m (e.g., the number of drones), C (e.g., the upper bound of the fire intensity), η (e.g., the perturbation parameter) as input. We vary these parameters and measure their impact on the performance of the Firefly algorithm. While varying parameter values, we consider synthetic and realistic graph structures based on Erdos-Renyi random graphs and the actual grid zoning structures (mirror those of Figure 1), respectively. The synthetic experiment results are omitted due to space restrictions. The performance of the Firefly algorithm is measured with respect to the optimal solution (i.e., the optimal knapsack value if we know the fire intensity at every location) at each time step. For the sake of comparisons, we consider the random firefighter allocation (without any information of the fires) where we randomly allocate resources to each zone (so that we don't exceed the resource constraints) as a baseline case. The performance will be presented as the ratio between the given algorithm (Firefly or random) and the optimal solution. We refer the ratio to be either Firefly/Opt or random/Opt (or more generally Alg/Opt). We report the average ratio over the time-steps.

In the realistic fire fighting instances, we consider the 10×10 grid graph that is based on a region in Marin County, California (See Figure 1). The priorities for each zone $i = 1, ..., n$ at time t is $k_{i,t} = k_i$ where k_i is drawn from a Gaussian distribution with mean 0.6 and standard deviation 0.2. We then classify zones as moderate, high or very high using data based on the Fire and Resource Assessment Program as in Figure 1. Using this classification, we generate the fire intensity at each zone as follows: every round the fire at a region is 0 with probability 0.5 and is $p \times C$ with probability 0.5 where C is the upper bound of the fire intensity and p is a random variable drawn from the following distributions: (a) A uniform distribution with lower bound 0.1 and upper bound 0.2 for moderate regions; (b) A uniform distribution with lower bound 0.4 and upper bound 0.5 for high regions; (c) A uniform distribution with lower bound 0.8 and upper bound 0.9 for very high regions. Given the 10 x 10 grid graph structure, we consider $r \in \{10, 50, 100, 200\}$, $m \in \{1, 5, 10, 50\}$ and $C \in \{10, 50\}$. For each of the combinations of the parameters, we generate 10 problem instances. The reported ratio is the average over the 10 problem instances (with different zone priorities). We omitted reporting the (small) standard deviations for the ease of presentations. The η value is set to be 0.5 for all of the experiments. Finally, we let $T = 200$ to be the maximum number of time-steps.

Observation 1 - Varying the Number of Drones. Figure 2 shows the experimental results when we vary the number of drones (m) fixing each combination of parameters. Each of the bright-colored solid lines and dark-colored dotted lines corresponds to the average ratio of Firefly/OPT and random/OPT over 200 time-steps, respectively. From the observation, regardless of r and C, Firefly/OPT increases as we increase the number of drones. In particular, even with a smaller number of drones (e.g., 5-10), Firefly achieves 60-70% of the optimal solution. The observation suggests that our algorithm will be benefited from having more drones, and drones are usually tools for providing meaningful information to fight fires.

Observation 2 - Varying the Number of Firefighters. Figure 3 shows the experimental results when we vary the number of firefighters (r). We observe that Firefly's performance depends on both m and C. In Figure 3 (a)-(c), Firefly/Opt increases when $C = 10$ and decreases when $C = 50$.

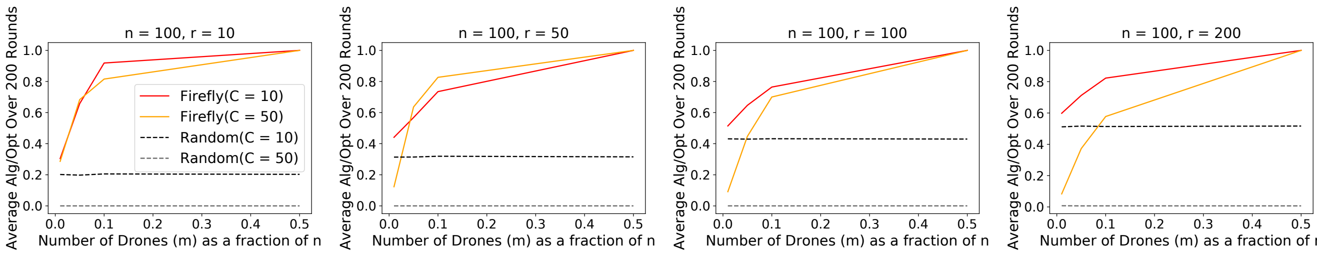

Figure 2: Realistic Instances Varying m: (a) r = 10, (b) r = 50, (c) r = 100, and (d) r = 200. The x-axis represents the number of drones, and the y-axis represents the average Alg/Opt over 200 time steps.

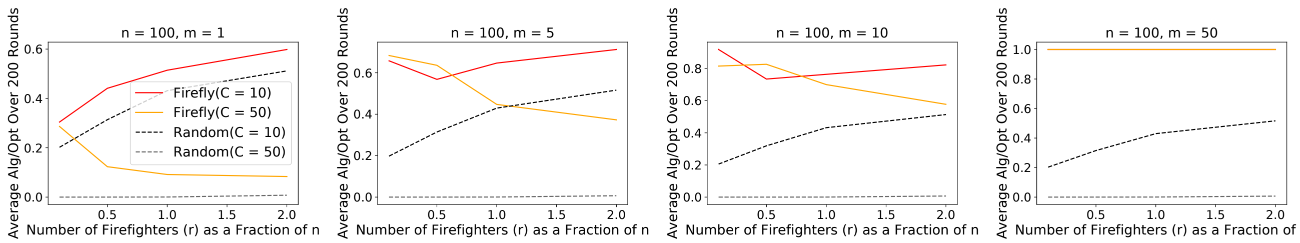

Figure 3: Realistic Instances Varying r: (a) m = 1, (b) m = 5, (c) m = 10, and (d) r = 50. The x-axis represents the number of firefighters, and the y-axis represents the average Alg/Opt over 200 time steps.

This could be due to the fact that the fire intensity is in a larger range with higher C and the optimal solution is harder to approximate. Figure 3 (d) shows that when we have a large enough number of drones ($m = 50$), Firefly could observe the whole graph and allocate firefighters optimally.

5 Conclusions

In this paper, we have proposed Firefly, a novel resource allocation problem to mitigate wildfire. In particular, our algorithm combines solutions for maximum set cover and knapsack problems to efficiently allocate both drones and firefighter units. Besides providing theoretical performance guarantees, we have also numerically evaluated the performance of our algorithm and demonstrated that it can achieve up to $80 - 90\%$ of the optimal performance even with a relatively small number of drones.

References

[Andrews, 2014] Patricia Andrews. Current status and future needs of the behaveplus fire modeling system. *International Journal of Wildland Fire*, 23, 2014.

[Belval *et al.*, 2014] Erin Belval, Yu Wei, and Michael Bevers. A mixed integer program to model spatial wildfire behavior and suppression placement decisions. *Canadian Journal of Forest Research*, 45, 07 2014.

[Bendix and Commons, 2017] Jacob Bendix and Michael Commons. Distribution and frequency of wildfire in california riparian ecosystems. *Environmental Research Letters*, 12, 05 2017.

[Boulton *et al.*, 2016] Chris Boulton, Humphfrey Shotton, and Hywel Williams. Using social media to detect and locate wildfires. 04 2016.

[Chmaj and Selvaraj, 2015] Grzegorz Chmaj and Henry Selvaraj. Distributed processing applications for uav/drones: a survey. In *Progress in Systems Engineering*, pages 449–454. Springer, 2015.

[Cruz *et al.*, 2008] Miguel Cruz, Martin Alexander, and Paulo Fernandes. Development of a model system to predict wildfire behaviour in pine plantations. *Australian Forestry*, 71:113–121, 06 2008.

[de Rigo *et al.*, 2013] Daniele de Rigo, Dario Rodriguez Aseretto, Claudio Bosco, Margherita Di Leo, and J. San-Miguel-Ayanz. *An architecture for adaptive robust modelling of wildfire behaviour under deep uncertainty*, volume 413, pages 367–380. 10 2013.

[Donovan and Rideout, 2003] Geoffrey H Donovan and Douglas B. Rideout. An integer programming model to optimize resource allocation for wildfire containment. 2003.

[Hochbaum and Pathria, 1998] Dorit S. Hochbaum and Anu Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)*, 45(6):615–627, 1998.

[HomChaudhuri and Cohen, 2010] B. HomChaudhuri and Kelly Cohen. Optimal fireline generation for wildfire fighting in uncertain and heterogeneous environment. pages 5638 – 5643, 08 2010.

[Hu and Ntaimo, 2009] Xiaolin Hu and Lewis Ntaimo. Integrated simulation and optimization for wildfire containment. *ACM Trans. Model. Comput. Simul.*, 19(4):19:1–19:29, November 2009.

[Keeley, 2009] Jon Keeley. Fire intensity, fire severity and burn severity: A brief review and suggested usage. *International Journal of Wildland Fire*, 18:116–126, 2009.

[Kyzirakos *et al.*, 2014] Kostis Kyzirakos, Manos Karpathiotakis, George Garbis, Charalampos Nikolaou, Konstantina Bereta, Ioannis Papoutsis, T. Herekakis, Dimitrios Michail, Manolis Koubarakis, and Kontoes Charalabos. Wildfire monitoring using satellite images, ontologies and linked geospatial data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 24, 01 2014.

[Lee *et al.*, 2012] Yo-Han Lee, Jeremy Fried, Heidi Albers, and Robert Haight. Deploying initial attack resources for wildfire suppression: Spatial coordination, budget constraints, and capacity constraints. *Canadian Journal of Forest Research*, 43:56–65, 11 2012.

[Malamud *et al.*, 2005] Bruce D. Malamud, James D. A. Millington, and George L. W. Perry. Characterizing wildfire regimes in the united states. 102(13):4694–4699, 2005.

[Neu and Bartók, 2013] Gergely Neu and Gábor Bartók. An efficient algorithm for learning with semi-bandit feedback. *ArXiv*, abs/1305.2732, 2013.

[Ntaimo *et al.*, 2012] Lewis Ntaimo, Julián Arrubla, Curt Stripling, Joshua Young, and Thomas Spencer. A stochastic programming standard response model for wildfire initial attack planning. *Canadian Journal of Forest Research*, 42:987–1001, 05 2012.

[Piñol *et al.*, 1998] Josep Piñol, J. M. López Terradas, and Francisco Lloret. Climate warming, wildfire hazard, and wildfire occurrence in coastal eastern spain. *Climatic Change*, 38:345–357, 1998.

[Pisinger, 1999] David Pisinger. Linear time algorithms for knapsack problems with bounded weights. *Journal of Algorithms*, 33(1):1 – 14, 1999.

[Rodriguez Aseretto *et al.*, 2013] Dario Rodriguez Aseretto, Daniele de Rigo, Margherita Di Leo, Ana Cortés, and J. San-Miguel-Ayanz. A data-driven model for large wildfire behaviour prediction in europe. *Procedia Computer Science (ISSN 1877-0509)*, 18:1861–1870, 06 2013.

[Sabit *et al.*, 2011] Hakilo Sabit, Adnan Al-Anbuky, and Hamid GholamHosseini. Wireless sensor network based wildfire hazard prediction system modeling. *Procedia Computer Science*, 5:106 – 114, 2011.

[Salis *et al.*, 2012] Michele Salis, Bachisio Arca, Alan Ager, Cinzia Fois, V. Bacciu, Pierpaolo Duce, and Donatella Spano. Extreme wildfire spread and behaviour: Case studies from north sardinia, italy. pages 10052–, 04 2012.

[Tsang *et al.*, 2013] Alan Tsang, Kate Larson, and Rob McAlpine. Resource sharing for control of wildland fires. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI'13, pages 1355–1361. AAAI Press, 2013.

[Westerling *et al.*, 2006] A. L. Westerling, H. G. Hidalgo, D. R. Cayan, and T. W. Swetnam. Warming and earlier spring increase western u.s. forest wildfire activity. 313(5789):940–943, 2006.

[Yanjun Li *et al.*, 2006] Yanjun Li, Zhi Wang, and Yeqiong Song. Wireless sensor network design for wildfire monitoring. In *2006 6th World Congress on Intelligent Control and Automation*, volume 1, pages 109–113, 2006.

[Yuan *et al.*, 2015] Chi Yuan, Youmin Zhang, and Zhixiang Liu. A survey on technologies for automatic forest fire monitoring, detection, and fighting using unmanned aerial vehicles and remote sensing techniques. *Canadian journal of forest research*, 45(7):783–792, 2015.