2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Wildfire Containment Using Multi-drone Systems: A Predict-Then-Optimize Approach

Aoran Cheng[1,†], Shijie Pan[1,2,†], Yiqi Sun[1], Kai Kang[3], Cristobal Pais[4], Yulun Zhou[5,*], and Zuo-Jun Max Shen[1,*]

***Abstract*— A multi-drone system coupled with data intelligence can be the future of wildfire fighting. However, multi-drone firefighting faces considerable challenges, such as complex environmental conditions, rapidly changing wildfire spread, and computational complexity in multi-drone operations. We address these challenges by developing a predict-then-optimize approach to enable multi-drone firefighting. First, we create a wildfire spread prediction model Firefighter-DResNet based on previous wildfire data. Next, we propose a Mixed-Integer Programming (MIP) model coupled with Dynamic Programming (DP) to facilitate efficient multi-drone task planning. We further use Chance-constrained Robust Optimization (CCRO) to ensure robust performance under varying conditions. After training with 75 varied wildfire environments, our approach demonstrates exceptional performance in controlling wildfires through multi-drone coordination. Our prediction model Firefighter-DResNet outperforms state-of-the-art models by 78.1% (20×20 environment) and 86.4% (40×40 environment) in terms of Intersection over Union (IoU), and achieves more than 54.7% and 27.3% higher precision. Subsequently, our chance-constrained optimizer succeeds in large-scale scenarios, achieving a 30% success rate in N=16 40×40 scenarios where all other baselines fail. This integrated solution significantly enhances the robustness and efficiency of autonomous multi-drone firefighting, marking a substantial leap towards practical deployment.**

I. Introduction

Wildfires have become more frequent and destructive in recent years, posing serious threats to life, property, and the environment. In 2025, the Los Angeles wildfire burned over 40,588 acres, nearly 90% of the size of Washington, D.C., significantly affecting Los Angeles [1]. In 2024, the United States faced approximately 36,000 wildfires, burning a total of over 7 million acres [2]. In China, 258 forest fires were reported, which resulted in 13 deaths and financial losses exceeding 2.68 billion yuan [3]. These devastating events underline the urgent need for more effective wildfire management strategies.

Multi-drone systems, as illustrated in Figure 1, are essential for enhancing firefighting efficiency. Their ability to precisely target fire points and execute effective tasks makes them invaluable tools in wildfire response operations. Most related studies have concentrated on deploying drones equipped with thermal infrared cameras for fire monitoring purposes. For instance, Shapira and Agmon [4] investigate optimal drone formations from a safety perspective, while Pham et al. [5] and Bailon-Ruiz et al. [6] propose more flexible drone allocation strategies. The former aims to maximize coverage of affected areas through 3D modeling, whereas the latter employs a Variable Neighborhood method to track fire fronts [7].

Fig. 1. Multi-drones system and fire base stations working together in firefighting (left) and a drone deploying firefighting bombs (right).

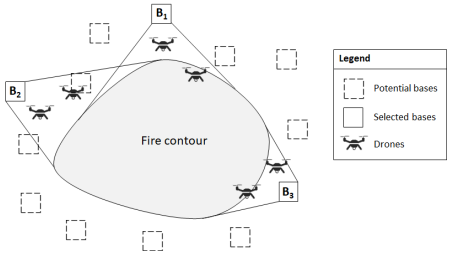

Fig. 2. A conceptual diagram of multi-drones system firefighting.

The project is supported by National Key RD Program of China (2022YFB3903704) and HKU Seed Fund for Basic Research (109001262).
†These authors contributed equally to this work. *Corresponding authors: Yulun Zhou and Zuo-Jun Max Shen. [1]A. Cheng, Y. Sun, and Z.-J. M. Shen are with the Faculty of Engineering, The University of Hong Kong, China. acheng39@connect.hku.hk; yiqisun@hku.hk; maxshen@hku.hk. [2]S. Pan is with the Dept. of ECE, Johns Hopkins University, USA. span34@jh.edu. [3]K. Kang is with the Faculty of Economics, Shenzhen University, China. kangkai@szu.edu.cn. [4]C. Pais is with Amazon Research, Seattle, WA, USA. crispais@amazon.com. [5]Y. Zhou is with the School of Innovation, The University of Hong Kong, China. yulunzhou@hku.hk.

Previous research has demonstrated that drone swarms can employ stigmergy-based coordination to cooperatively detect and extinguish fire spots [8], and that unmanned aerial vehicles can collaborate with manned aircraft for wildfire surveillance and suppression [9]. Unlike existing studies, our work extends beyond fire monitoring by incorporating the use of drones to extinguish fire points (as illustrated in Figure 2), thereby directly influencing fire progression. Rather than relying on rule-based strategies, we formulate the problem as an optimization model with theoretical guarantees. By accounting for drone capacity constraints and the transit process between fire points and bases, our model integrates these critical factors into task allocation, enhancing coordination and adaptability in dynamic wildfire scenarios [10].

Multi-drone firefighting faces significant challenges, particularly in allocation and coordination. The complexity of wildfire environments often results in suboptimal solutions and inefficient resource allocation [11]. Existing wildfire prediction models, such as cellular automata [12] and Markov chains [13], struggle to integrate with optimization strategies, thereby further hindering operational efficiency. However, recent advancements in prediction and optimization have demonstrated considerable promise. Deep learning-based models such as Swin-Unet [14] effectively capture spatiotemporal fire dynamics, significantly improving prediction accuracy.Additionally, PatchwiseConvLSTM [15] captures the spatiotemporal dynamics of wildfire spread by learning from localized spatial patterns over time, while Causal-GNN [16] models the causal relationships between environmental factors and fire progression. Together, these approaches enhance the capability to model wildfire dynamics, providing enhanced interpretability and enabling more precise risk assessments. Meanwhile, MA-Net [17] integrates multimodal data sources, improving predictive accuracy and outperforming automata modeling approaches.

In drone planning, Genetic Algorithm for Routing and Scheduling in Wildfire Suppression (GARST) [18] have enhanced the coordination and scheduling of drones for wildfire suppression. These approaches address the challenges posed by dynamic fire environments by enabling proactive and adaptive decision-making processes. Resource-Efficient Decentralized Sequential (REDS) [19] improves drone efficiency by optimizing resource allocation in real-time, while Ant Colony Initialized Genetic Algorithm (ACIGA) [20] provides balanced path planning to ensure continuous monitoring, tracking, and suppression of wildfires. Although these methods have improved drone coordination efficiency to some extent, they still encounter certain limitations in practical applications. For instance, Criss-Cross Polar Light Optimization (CCPLO) [21] enhances coordination in multi-drone planning; however, its adaptability to dynamic and complex environments remains limited, particularly in wildfire scenarios characterized by high variability and demanding real-time feedback requirements, where it may fail to provide efficient resource allocation and optimal path planning.

In this research, we develop a data-driven wildfire prediction model, **Firefighter-DResNet**, to guide optimal drone assignments. While CNN-based methods are commonly used in fire prediction, our model leverages a deep residual block architecture, improving feature extraction and learning of complex wildfire patterns. By integrating historical data, Firefighter-DResNet more accurately simulates fire spread dynamics. Compared to traditional grid-based methods [22], it better handles multidimensional complexities and supports efficient multi-drone deployment. To optimize drone resource allocation, we incorporate advanced techniques such as Mixed-Integer Programming (MIP), Dynamic Programming (DP), and Chance-Constrained Robust Optimization (CCRO). This approach ensures more effective resource allocation in complex, dynamic environments, maximizing long-term benefits and reducing computational burdens. As a result, our model enables drones to respond quickly and effectively during critical moments, improving wildfire management and emergency response.

II. Methods

Figure 3 illustrates the predict-then-optimize approach. First, wildfire spread dynamics are predicted by Firefighter-DResNet. The resulting predictions, $\boldsymbol{x}_t$, are then input into the optimization phase (2) to solve the MIP problem.

A. Wildfire prediction using Firefighter-DResNet

The training data is generated using the Cell2Fire model [23], with a dataset that includes 75 wildfire environments, such as Grassland, Mica Creek, and Revelstoke Forests. These environments are parameterized using realistic fuel, weather, moisture, and topography conditions, and are widely used as benchmark scenarios in the wildfire modeling community. Each sample consists of the fire state map at time t, $\mathbf{Map}_t$, and the corresponding drone intervention map, $\boldsymbol{x}_t$, where drone interventions convert burning pixels into extinguished ones. Given the fire map at time t, $\mathbf{Map}_t$, we assume that the drone intervention extinguishes the fire with a certain probability, while the remaining fire continues to spread. The fire state at the next time step, $\mathbf{Map}_{t+1}$, is simulated using the modified map $\mathbf{Map}_t^{'}$ and the Cell2Fire model. Data augmentation is performed using geometric transformations, increasing the number of 20×20 environment pairs to 3,120 and the number of 40×40 environment pairs to 2,000. To mitigate overfitting to simulation artifacts, training and test environments are strictly separated, and we include diverse

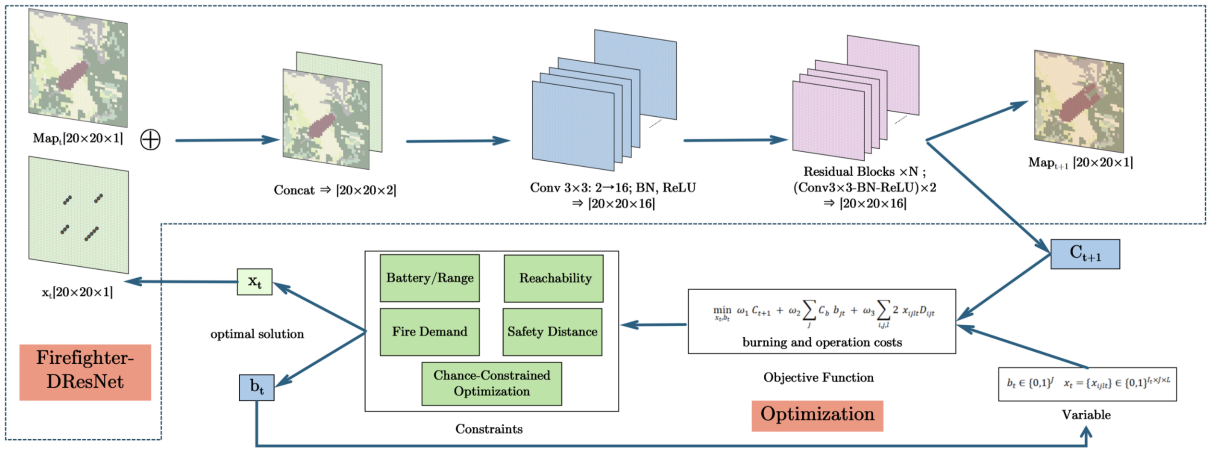

Fig. 3. Predict-then-optimize approach flowchart.

landscapes to promote generalization. All reported results are consistent across multiple trials, with negligible variation across runs. For fairness, baseline models are either adopted from official implementations or re-implemented and tuned under the same training settings as our method. To accurately predict wildfire spread, we develop the Firefighter-DResNet model, which takes the fire state map $\mathbf{Map}_t$ and drone intervention map $\boldsymbol{x}_t$ as inputs. The fire state map represents the wildfire's state at time t, with matrix elements having values of 0 (no fire), 1 (fire burning), or 2 (fire extinguished). The drone intervention map represents the drone actions, and the two maps are concatenated and fed into the network for prediction. The Firefighter-DResNet uses deep residual learning to model complex wildfire propagation dynamics.

- **Convolutional Layer and Deep Residual Block:** In the convolutional layer, we use 16 convolutional kernels of size 3×3 to generate feature maps of dimensions $20 \times 20 \times 16$ (or $40 \times 40 \times 16$), followed by batch normalization and ReLU activation to improve training performance. In addition,the deep residual blocks, consists of two stacked 3×3 convolutional layers, also equipped with batch normalization and ReLU activation; the design of the residual connections effectively mitigates the vanishing gradient problem, enabling the network to learn more complex wildfire spread patterns.
- **Prediction Layer and Loss Function:** The final output of the model is the predicted fire state map $\mathbf{Map}_{t+1}$. The training goal is to minimize the error between the predicted and actual fire state. We use the BCEWithLogitsLoss (binary cross-entropy with logits) loss function, which accounts for both fire prediction accuracy and burning costs, thereby optimizing the model's performance in complex environments. This loss function combines sigmoid activation with binary cross-entropy, improving stability and effectiveness for training binary classification models.

B. Optimal Multi-drones Task Allocation

Figure 4 illustrates an example of optimized drone flight paths in a 20x20 forest environment covered by four drones. The four drones take off from four base stations located at four corners of the 20x20 forest and fly towards optimized targets collectively to minimize burning costs over four fire spread periods. The red points in Figure 4 illustrate current fire points, while the blue points illustrate predicted fire spread points in the next period as predicted by the Firefighter-DResNet.

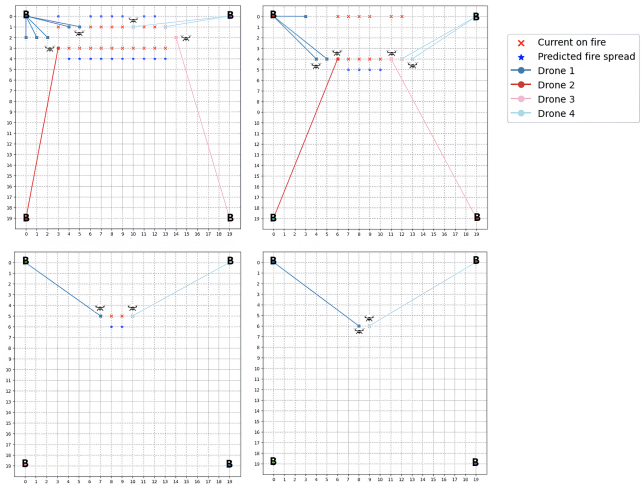

Fig. 4. Optimized drone flight paths from our method in a 20x20 forest environment covered by four drones.

We treat the base activation decisions and drone task allocation as a unified optimization problem, with the chance constraint (7) converted into a robust equivalent form to ensure solvability. The objective function combines the costs from forest burning and multi-drone deployment, including base activation and drone flying costs, each weighted by appropriate factors. Our decision variables include both base activation decisions, b_{jt}, and drone task allocation decisions, x_{ijlt}, where b_{jt} represents the activation state of base j at time t, and x_{ijlt} indicates the decision of drone l departing from base j to extinguish the fire at point i at time t, and eventually returning to the base.

To simplify the formulation and calculation, we make the following assumptions. Each drone has a capacity of 1 unit and speed α. Battery consumption and flight time are proportional to the total distance flown, with a coefficient of D_l. Batteries are replaced at the end of each time slot when the low capacity condition $\{u_{lt} \leq s\}$ is met. This allows each drone to perform multiple round trips (forward $\rightarrow$ release firefighting bomb $\rightarrow$ backward) within each time period of length Δ and return to its base. In the firefighting task, we assume that a drone has a 0.8 probability of extinguishing a fire point [24]. We also perform sensitivity analysis for probabilities of 0.7 and 0.9 to assess how these variations affect task allocation and overall cost. To account for uncertainty in the firefighting bomb delivery time, we assume it is influenced by factors such as heavy smoke or terrain. The bomb delivery time at each grid point, $\boldsymbol{\tau} = [\tau_1, \tau_2, ..., \tau_I]$, follows a joint distribution $\mathbb{P}_I$ with a positive definite covariance matrix, though the specific distribution form is unknown [25]. Additionally, we define a set $[I_t]$ containing all active fire points at time t, with intensity $\boldsymbol{g}_t$. The size of I_t is much smaller than I, the total number of grid points. In this set $[I_t]$, only fire points are considered for drone intervention, significantly reducing the feasible region of decision variables.

a) Mixed-Integer Programming: We develop a MIP model to optimize task allocation and route planning for drones in a real-time, dynamic firefighting environment. The optimization model is solved using the PuLP solver in Python3. The decision variables in optimization use binary values to represent the activation of each base j, b_{jt}. Besides, the assignment of drone tasks to drone l, which departs from base j to extinguish the fire of point i at time t (and finally returns), is defined by the binary decision variables x_{ijlt}. There are totally L drones in all J bases. In equation (2)-(3), $\mathbf{Map}_t$ is the fire intensity map at the start of time slot t, which includes the fire point set $[I_t]$ with intensity $\boldsymbol{g}_t$. The set $[I_t]$ represents the active fire points at time t, while $\boldsymbol{g}_t$ is the intensity of the fire at each point, which directly influences the urgency and task allocation. (Normally, $\boldsymbol{g}_t$ equals to one.) $\boldsymbol{x}_t$ is a vector of task assignments x_{ijlt}, representing the drone's specific task allocation for each fire point and period t. C_{t+1} is the future burning cost at the start of the next time slot $t+1$, i.e., the cost incurred due to the continuous burning of uncontrolled fire points in the next time slot as a result of the resource allocation decision in the previous time slot, denoted as $O_c[f]$. C_b is the cost of single base activation.

Our model focuses on three key objectives: minimizing the burning cost in the next period, optimizing drone battery usage, and task route efficiency, including base activation operations to ensure rapid response and wildfire threat mitigation. These objectives are combined using predefined weights for each objective. For example, ω_1 assigns higher importance to minimizing the burning cost, while ω_2 and ω_3 focus on optimizing base activation decision and drone task routing for efficiency.

$$\min_{\boldsymbol{x}_t,\, \boldsymbol{b}_t} \omega_1 C_{t+1} + \omega_2 \sum_j C_b b_{jt} + \omega_3 \sum_{i,j,l} 2 x_{ijlt} D_{ijt} \tag{1}$$

$$\text{s.t. } G(\mathbf{Map}_t,\, \{\boldsymbol{x}_{ijlt}\}) = \{\mathbf{Map}_{t+1}\} \tag{2}$$

$$\sum_f \{\mathbf{Map}_{t+1}[f] = 1 \wedge \mathbf{Map}_t[f] = 0\} * O_c[f] = C_{t+1} \tag{3}$$

$$\sum_i 2 x_{ijlt} D_{ijt} \leq D_l u_{lt},\ j = 1, ..., J,\ l = 1, ..., L \tag{4}$$

$$x_{ijlt} \leq \min\{b_{jt},\, y_{jl},\, \pi_{ijt}\},\ i = 1, ..., I_t,\ j = 1, ..., J,\ l = 1, ..., L \tag{5}$$

$$\sum_{l,j} x_{ijlt} \leq g_{it},\ i = 1, ..., I_t \tag{6}$$

$$\mathcal{P}_{\mathbb{P}_I}\{\sum_{i,j} 2 x_{ijlt} (D_{ijt}\alpha^{-1} + \frac{1}{2}\tau_i) \leq \triangle\} \geq 1 - \delta,\ \forall \mathbb{P}_I \in \mathbb{F}_I,\ l = 1, ..., L \tag{7}$$

$$\pi_{ijt} b_{jt} D_{safe} \leq D_{ijt},\ i = 1, ..., I_t,\ j = 1, ..., J \tag{8}$$

$$x_{ijlt} \in \{0, 1\};\ b_{jt} \in \{0, 1\}$$

The model inputs include the current and predicted locations of fire points ($\mathbf{Map}_t$ and $\mathbf{Map}_{t+1}$), the fire point-base allocation π_{ijt} (which represents the assignment of fire points to base stations), and drone-specific parameters such as battery capacity u_{lt}, distance between bases and fire points D_{ijt}, and the drone-base affiliation y_{jl}. These parameters affect task allocation by considering environmental factors, drone capabilities, and the urgency of the fire's spread.

The MIP model incorporates several constraints, such as battery limitations constraint (4), fire-base-drone affiliation availability constraint (5), firefighting time constraint (6), limited operating time constraint (a chance-constraint ensuring the task can be completed within Δ time with at least $1 - \delta$ probability, (7)), and the safe distance constraint

D_{safe} between the base and the fire's contour (8). These constraints ensure that the model produces feasible and safe task schedules while optimizing for the objectives.

b) Myopic Dynamic Programming: We adapt a Myopic Dynamic Programming (Myopic DP) framework, extending the single-period MIP model to optimize drone deployment across multiple time slots. This approach considers only the costs and task allocations for the current and next stages, addressing the long-term unpredictability of fire in wildfire environments. Since prediction models perform poorly when forecasting over longer periods, the myopic setting helps improve accuracy, ensuring safer and more efficient firefighting operations [17]. The transition constraints within the DP framework equations (9)-(11) are as follows:

$$v_{l(t+1)} = \lceil \frac{1}{M} \max\{0,\ u_{lt} - \frac{1}{D_l} \sum_{i,j} 2x_{ijlt} D_{ijt} - s\} \rceil,$$
$$l = 1, ..., L,\ t = 0, ..., T-1 \tag{9}$$

$$u_{l(t+1)} = v_{l(t+1)}(u_{lt} - \frac{1}{D_l} \sum_{i,j} 2x_{ijlt} D_{ijt}) + (1 - v_{l(t+1)}),$$
$$l = 1, ..., L,\ t = 0, ..., T-1 \tag{10}$$

$$g_{iT} = \sum_{j,l} x_{ijlT},\ i = 1, 2, ..., I_T \tag{11}$$

where equations (9) and (10) ensure each drone has at least s energy at the start of each time slot, indicated by u_{lt}, and prepares them for subsequent tasks. M is a large positive penalty number to regularize the remaining battery level to $[0, 1]$. Equation (11) defines the terminal state constraint, implying that all fire points must be extinguished within T periods. We should ensure all fire points are extinguished at the end of period T. MIP integration optimizes drone operations for efficient resource use and timely firefighting, while Myopic Dynamic Programming enables real-time management of drones in complex wildfire environments, improving response speed and operational efficiency.

*c) **Chance-constraint Robust Optimization**:* In our MIP, we have the chance-constraint (7) as follows,

$$\mathcal{P}_{\mathbb{P}_I}\{\sum_{i,j} 2x_{ijlt}(D_{ijt}\alpha^{-1} + \frac{1}{2}\tau_i) \leq \triangle\} \geq 1 - \delta,$$
$$\forall\, \mathbb{P}_I \in \mathbb{F}_I,\ l = 1, ..., L,$$

where $\mathbb{F}_I$ is a Mean-Variance (MV) robust set given by:

$$\mathbb{F}_I = \left\{\mathbb{P}_I \in \mathcal{P}_0((R)^I) |\ \mathbb{E}_{\mathbb{P}_I}[\boldsymbol{\tau}] = \boldsymbol{\mu},\ \mathrm{Var}_{\mathbb{P}_I}[\boldsymbol{\tau}] = \boldsymbol{\Sigma}\right\} \tag{12}$$

The uncertainty set accounts for the bomb-delivery time uncertainty at each grid point on the map. To simplify computational complexity, we focus on the uncertainty set $\mathbb{P}_{I_t}$ and $\mathbb{F}_{I_t}$ for the I_t fire points in each time slot t. The robust chance-constraint initially leads to an intractable optimization problem with different metrics. However, inspired by the work of Last-Mile Delivery [26], we derive an equivalent constraint formulation, combining a second-order positive definite constraint (13) and a linear constraint (14) as shown in Proposition 1. This formulation allows for an efficient solution.

Proposition 1. *In time slot t, the Chance-constraint (7) in MIP is equivalent to the following constraints,*

$$4\boldsymbol{x}^\top(\frac{1-\delta}{4\delta} A^\top \boldsymbol{\Sigma} A - (\alpha^{-1}\boldsymbol{D} + \frac{1}{2}A^\top\boldsymbol{\mu})(\alpha^{-1}\boldsymbol{D}$$
$$+ \frac{1}{2}A^\top\boldsymbol{\mu})^\top)\boldsymbol{x} + 4\Delta\boldsymbol{x}^\top(\alpha^{-1}\boldsymbol{D} + \frac{1}{2}A^\top\boldsymbol{\mu}) - \Delta^2 \leq 0,$$
$$l = 1, ..., L \tag{13}$$

$$2\alpha^{-1}\boldsymbol{x}^\top\boldsymbol{D} - \Delta + \boldsymbol{x}^\top A^\top\boldsymbol{\mu} \leq 0,\ l = 1, ..., L \tag{14}$$

where $\boldsymbol{x}$ and $\boldsymbol{D}$, are vectors with dimension I_tJ, for decision variable x_{ijlt}, and distance between bases and fire points D_{ijt}, respectively. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are defined in equation (12). The matrix A is a pre-conditioner to compress the dimension of the vector from I_tJ to I_t. Specifically, $A = \begin{bmatrix} \mathbf{1} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{1} \end{bmatrix}_{I_t \times I_t J}$ *,where* $\mathbf{1} = [1, 1, ..., 1]_J$ *and* $\mathbf{0} = [0, 0, ..., 0]_J$.

Proposition 1 is proved using the result from Ghaoui, Oks, and Oustry [27]. By shifting the robust set $\mathbb{F}_I$ to zero mean and squaring both sides of the original robust equivalent form, we obtain equation (13). Equation (14) eliminates the feasibility region introduced by the squaring operation. For a sufficiently small δ (high confidence), Equation (13) is always positive definite if $\boldsymbol{\Sigma}$ is positive definite.

III. Data and Experiments

A. Wildfire Spread Prediction Performance

Table I presents a comprehensive performance comparison of several wildfire spread prediction models, including Swin-UNet, MA-Net, PatchwiseConvNN, Causal-GNN, CNN, and Firefighter-DResNet, evaluated on 20×20 and 40×40 forest environments. Notably, the comparison between CNN and Firefighter-DResNet also serves as an ablation study on the effect of deep residual blocks, since the two models share the same convolutional backbone but differ in residual learning design. The evaluation metrics include accuracy, precision, recall, F1 score, and Intersection over Union (IoU).

In terms of precision, Firefighter-DResNet achieves 0.8184 for the 20×20 environment and 0.6443 for the 40×40 environment, surpassing CNN by 13.1% and 47.7%, respectively. Compared to MA-Net, the improvement is 54.7% for 20×20 and 27.3% for 40×40. Higher precision reduces

false positives, which helps avoid mistakenly identifying non-fire areas and enhances the efficiency of firefighting resource allocation. For recall, Firefighter-DResNet achieves 0.8468 for 20×20 and 0.7933 for 40×40, outperforming MA-Net by 104.7% and 90.8%, respectively. Higher recall ensures that true wildfire regions are identified, minimizing false negatives and ensuring a more comprehensive wildfire response.

The F1 score, balancing both precision and recall, also demonstrates the superiority of Firefighter-DResNet. The model achieves F1 scores of 0.8322 for 20×20 and 0.7111 for 40×40, which are 79.3% and 55.7% higher than MA-Net, respectively. This indicates that Firefighter-DResNet excels at balancing false positives and false negatives, providing more reliable predictions. In IoU, Firefighter-DResNet also leads with values of 0.5384 for 20×20 and 0.5517 for 40×40, improving by 78.1% and 86.4% over MA-Net. Higher IoU values indicate better delineation of wildfire boundaries, which is crucial for accurately predicting fire spread.

Although Firefighter-DResNet shows slightly lower accuracy in the 20×20 environment, its significant advantages in other metrics demonstrate its robustness, making it particularly suitable for complex wildfire prediction environments.

B. Optimization with Chance-Constrained Robust Allocation

Tables II and III present a comparison of task-allocation performance and success rates between CCRO and a variety of baselines in 20×20 and 40×40 grid scenarios. The baselines cover both classical optimization methods (e.g., Box Uncertainty Robust Optimization—an ablation that excludes chance constraints), Genetic Algorithm is widely used in recent multi-drone literature. All methods are evaluated over 10 independent runs and reported in terms of success rate and average cost, assuming a single-drone extinguishing probability of $p_{\text{success}} = 0.8$. A task is considered successful if the fire is extinguished within 5 rounds; if the fire is not extinguished within this time frame, it is counted as a failure.

CCRO demonstrates strong performance across both grid sizes. In the 20×20 grid (Table II), as N increases from 4 to 12, CCRO achieves success rates of 30%, 90%, and 100%, with average costs decreasing from 142.14 to 72.94, indicating efficient resource utilization and scalability. In the 40×40 grid (Table III), CCRO attains success rates of 10%, 30%, and 40% for $N = 12, 16, 20$, respectively, with costs of 272.48, 294.93, and 186.41. Notably, at $N = 12, 16$, CCRO is the only algorithm to achieve a non-zero success rate (40%).

Box Uncertainty Robust Optimization, which excludes chance constraints, underperforms CCRO in both grids. In 20×20, it reaches 80% and 100% success at $N = 8$ and 12, but with higher costs than CCRO (e.g., 146.83 vs. 113.85 at $N = 8$). In 40×40, its success rate drops to 0% at $N = 12, 16, 20$.It demonstrates the superiority of the CCRO Method. This limited scalability highlights the value of chance constraints in handling uncertainty. Resource Efficient Decentralized Sequential (REDS), a heuristic auction-based method, performs poorly overall. In 20×20, it only succeeds at $N = 12$ with 40% success and a cost of 125.63, which is 72% higher than CCRO's cost at the same task size. In 40×40, REDS achieves 10% success at $N = 20$, with costs of 195.42. While its cost at $N = 16$ is lower than CCRO's, its success rate is only one-fifth of CCRO's (10% vs. 50%). At $N = 20$, CCRO achieves four times the success rate (40% vs. 10%) compared to REDS.

Table IV presents a sensitivity analysis of CCRO, Box Uncertainty Robust Optimization, and REDS under varying extinguishing probabilities ($p_{\text{success}} = 0.9, 0.8, 0.7$). CCRO consistently outperforms the other methods across all settings. At $p_{\text{success}} = 0.9$, CCRO achieves 100% success at $N = 12$ in 20×20, and 50% at $N = 20$ in 40×40, while Box Uncertainty reaches 100% only at $N = 12$ in 20×20 and unable to complete any rescues in 40×40. As p_{success} decreases to 0.7, CCRO maintains 70% success at $N = 12$ in 20×20 and 20% at $N = 20$ in 40×40. Actually, at $p = 0.7$, Box Uncertainty in 40×40 is 0% for all N. REDS only achieves non-zero success at $p = 0.9$, 0.8 for $N = 12$ in 20×20 (70% and 40%) and at $N = 20$ in 40×40 (10% at $p = 0.9, 0.8$), but fails completely at $p = 0.7$ in 40×40.

The superior performance of CCRO stems from its explicit incorporation of chance constraints, which allows it to account for uncertainty in fire extinguishing outcomes. In contrast, Box Uncertainty Robust Optimization assumes worst-case conditions and becomes overly conservative, leading to failure when actual success probabilities deviate from the assumed threshold. REDS, lacking probabilistic modeling, cannot adapt to uncertainty and thus exhibits consistently low success rates. These results underscore the importance of chance-constrained methods for robust and efficient multi-drone task allocation in wildfire response. In summary, CCRO outperforms all compared methods in terms of success rate, average cost, and scalability across varying grid sizes and extinguishing probabilities. Its ability to balance resource efficiency with robustness makes it a promising solution for real-world wildfire management.

IV. Conclusion & Future Work

This study presents an innovative predict-then-optimize approach, combining the Firefighter-DResNet wildfire prediction model with drone resources through chance-constrained robust optimization (MIP-CCRO) and dynamic programming (DP), which improves firefighting efficiency, reduces costs, and accounts for uncertainty in drone intervention success. While these advancements are significant,

TABLE I

PERFORMANCE COMPARISON OF WILDFIRE SPREAD PREDICTION MODELS ON 20×20 AND 40×40 GRIDS. RESULTS ARE AVERAGED OVER 5 RUNS WITH RANDOM SEEDS. EXPERIMENTS WERE RUN ON A 2.4 GHZ QUAD-CORE INTEL CORE I5.

Metric	**Swin-UNet**		**MA-Net**		**PatchwiseConvNN**		**GCN**		**CNN**		**Firefighter-DResNet**	
	20×20	40×40	20×20	40×40	20×20	40×40	20×20	40×40	20×20	40×40	20×20	40×40
Accuracy	0.9677	0.9743	**0.9791**	0.8993	0.9722	0.9690	0.8410	0.8623	0.7374	0.9771	0.8331	**0.9879**
Precision	0.2309	0.2669	0.5290	0.5062	0.3881	0.3055	0.3698	0.3494	0.7231	0.4367	**0.8184**	**0.6443**
Recall	0.2042	0.2165	0.4136	0.4159	0.4713	0.5002	0.6121	0.6591	0.7510	0.7426	**0.8468**	**0.7933**
F1 score	0.2168	0.2391	0.4642	0.4566	0.4257	0.3793	0.4610	0.4567	0.7372	0.5499	**0.8322**	**0.7111**
IoU	0.1216	0.1358	0.3023	0.2959	0.2704	0.2620	0.2995	0.2960	0.4492	0.3793	**0.5384**	**0.5517**

TABLE II

TASK ALLOCATION PERFORMANCE IN 20X20 GRID WITH $p_{\text{SUCCESS}} = 0.8$ UNDER DIFFERENT METHODS: PROPOSED VS. OTHER APPROACHES

Algorithm	**N=4**		**N=8**		**N=12**	
	Succ. (%)	**Avg. cost (succ.)**	**Succ. (%)**	**Avg. cost (succ.)**	**Succ. (%)**	**Avg. cost (succ.)**
Chance-Constraint Robust Optimization	**30%**	**142.14**	**90%**	**113.85**	**100%**	**72.94**
Box Uncertainty Robust Optimization	20%	142.11	80%	146.83	100%	83.00
Heuristic Greedy	0%	N/A	50%	131.52	90%	102.58
Ant Colony Initialized Genetic Algorithm	10%	129.27	60%	114.48	90%	87.65
Variable Neighborhood Search	0%	N/A	50%	129.74	90%	100.23
Criss-Cross Polar Lights Optimization	10%	143.56	50%	137.72	50%	108.91
Resource Efficient Decentralized Sequential	0%	N/A	0%	N/A	40%	125.63

TABLE III

TASK ALLOCATION PERFORMANCE IN 40X40 GRID WITH $p_{\text{SUCCESS}} = 0.8$ UNDER DIFFERENT METHODS: PROPOSED VS. OTHER APPROACHES

Algorithm	**N=12**		**N=16**		**N=20**	
	Succ. (%)	**Avg. cost (succ.)**	**Succ. (%)**	**Avg. cost (succ.)**	**Succ. (%)**	**Avg. cost (succ.)**
Chance-Constraint Robust Optimization	**10%**	**272.48**	**30%**	**294.93**	**40%**	**186.41**
Box Uncertainty Robust Optimization	0%	N/A	0%	N/A	0%	N/A
Heuristic Greedy	0%	N/A	0%	N/A	0%	N/A
Ant Colony Initialized Genetic Algorithm	0%	N/A	0%	N/A	0%	N/A
Variable Neighborhood Search	0%	N/A	0%	N/A	0%	N/A
Criss-Cross Polar Lights Optimization	0%	N/A	0%	N/A	0%	N/A
Resource Efficient Decentralized Sequential	0%	N/A	0%	N/A	10%	195.42

we acknowledge that external validation using real wildfire data remains a key direction for future research. Future work will compare myopic and non-myopic planning strategies to better quantify long-term performance trade-offs. Additionally, future research should address practical challenges in deploying drones for wildfire management, such as system safety, failure risks, and adaptability in complex environments. Integrating uncertainty-aware methods and multi-agent reinforcement learning will further optimize drone coordination and decision-making, enhancing the adaptability and efficiency of wildfire management systems.

REFERENCES

[1] Reuters, "California wildfire explodes, becomes largest in the us," 2025, accessed: February 10, 2025.

[2] Guancha, "Three massive wildfires are burning simultaneously in southern california, sending huge 'fire cumulus' clouds into the sky. international news report," 2024, accessed: January 10, 2025.

[3] Ministry of Emergency Management of the PRC, "Statistics from the ministry of emergency management," 2025, accessed: January 10, 2025.

[4] Y. Shapira and N. Agmon, "Path planning for optimizing survivability of multi-robot formation in adversarial environments," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 4544–4549.

[5] H. X. Pham, H. M. La, D. Feil-Seifer, and M. Deans, "A distributed control framework for a team of unmanned aerial vehicles for dynamic

TABLE IV

SENSITIVITY ANALYSIS: SUCCESS RATE OF CHANCE CONSTRAINT ROBUST OPTIMIZATION, BOX UNCERTAINTY ROBUST OPTIMIZATION, AND RESOURCE EFFICIENT DECENTRALIZED SEQUENTIAL UNDER VARYING FIRE EXTINGUISHING PROBABILITIES ON 20X20 AND 40X40 GRID.

Algorithm	p_{success}	**20×20 Grid**			**40×40 Grid**		
		N=4	**N=8**	**N=12**	**N=12**	**N=16**	**N=20**
Chance Constraint Robust Optimzation	0.9	40%	90%	100%	30%	50%	50%
	0.8	30%	90%	100%	10%	30%	40%
	0.7	20%	70%	70%	0%	10%	20%
Box Uncertainty Robust Optimization	0.9	30%	90%	100%	0%	20%	20%
	0.8	20%	80%	100%	0%	0%	0%
	0.7	10%	30%	80%	0%	0%	0%
Resource Efficient Decentralized Sequential	0.9	0%	0%	70%	0%	0%	10%
	0.8	0%	0%	40%	0%	0%	10%
	0.7	0%	0%	30%	0%	0%	0%

wildfire tracking," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 6648–6653.

[6] R. Bailon-Ruiz, S. Lacroix, and A. Bit-Monnot, "Planning to monitor wildfires with a fleet of uavs," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4729–4734.

[7] D. Lindenmayer, P. Zylstra, and M. Yebra, "Adaptive wildfire mitigation approaches," *Science*, vol. 377, no. 6611, pp. 1163–1164, 2022.

[8] I. L. H. Alsammak, M. A. Mahmoud, S. S. Gunasekaran, A. N. Ahmed, and M. AlKilabi, "Nature-inspired drone swarming for wildfires suppression considering distributed fire spots and energy consumption," *IEEE Access*, vol. 11, p. 50962–50983, 2023.

[9] M. Al-Husseini, K. H. Wray, and M. J. Kochenderfer, "Hierarchical framework for optimizing wildfire surveillance and suppression using human-autonomous teaming," *Journal of Aerospace Information Systems*, vol. 21, no. 10, p. 790–811, 2024.

[10] S.-H. Chung, "Applications of smart technologies in logistics and transport: A review," *Transportation Research Part E: Logistics and Transportation Review*, vol. 153, p. 102455, 2021.

[11] Z. Wang and J.-B. Sheu, "Vehicle routing problem with drones," *Transportation research part B: methodological*, vol. 122, pp. 350–364, 2019.

[12] A. Alexandridis, D. Vakalis, C. I. Siettos, and G. V. Bafas, "A cellular automata model for forest fire spread prediction: The case of the wildfire that swept through spetses island in 1990," *Applied Mathematics and Computation*, vol. 204, no. 1, pp. 191–201, 2008.

[13] D. Boychuk, W. J. Braun, R. J. Kulperger, Z. L. Krougly, and D. A. Stanford, "A stochastic forest fire growth model," *Environmental and Ecological Statistics*, vol. 16, pp. 133–151, 2009.

[14] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, and M. Wang, "Swin-unet: Unet-like pure transformer for medical image segmentation," in *Computer Vision – ECCV 2022 Workshops*. Springer Nature Switzerland, 2023, p. 205–218.

[15] A. Masrur, M. Yu, and A. Taylor, "Capturing and interpreting wildfire spread dynamics: attention-based spatiotemporal models using convlstm networks," *Ecological Informatics*, vol. 82, p. 102760, Sept. 2024.

[16] S. Zhao, I. Prapas, I. Karasante, Z. Xiong, I. Papoutsis, G. Camps-Valls, and X. X. Zhu, "Causal graph neural networks for wildfire danger prediction," 2024. [Online]. Available: https://arxiv.org/abs/2403.08414

[17] D. Shadrin, S. Illarionova, F. Gubanov, K. Evteeva, M. Mironenko, I. Levchunets, R. Belousov, and E. Burnaev, "Wildfire spreading prediction using multimodal data and deep neural network approach," *Scientific reports*, vol. 14, no. 1, p. 2606, 2024.

[18] J. John and S. Sundaram, "Genetic algorithm-based routing and scheduling for wildfire suppression using a team of uavs," in *2024 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2024, pp. 1–8.

[19] J. John, S. Velhal, and S. Sundaram, "A resource-efficient decentralized sequential planner for spatiotemporal wildfire mitigation," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 11 469–11 482, 2025.

[20] X. Zhan, Y. Chen, X. Chen, and W. Zhang, "Balanced multi-uav path planning for persistent monitoring," *Robotica*, vol. 43, p. 332–349, 2025.

[21] Z. Xu, Z. Wang, R. Liu, C. Huang, Y. Shi, M. Wang, and H. Chen, "Efficient multi-uav path planning in dynamic and complex environments using hybrid polar lights optimization," *Journal of King Saud University — Computer and Information Sciences*, vol. 37, no. 6, p. 125, 2025.

[22] P. Johnston, J. Kelso, and G. J. Milne, "Efficient simulation of wildfire spread on an irregular grid," *International Journal of Wildland Fire*, vol. 17, no. 5, pp. 614–627, 2008.

[23] C. Pais, J. Carrasco, D. L. Martell, A. Weintraub, and D. L. Woodruff, "Cell2fire: A cell-based forest fire growth model to support strategic landscape management planning," *Frontiers in Forests and Global Change*, vol. 4, 2021.

[24] Y. Wang, F. Gao, and M. Li, "Probabilistic path planning for uavs in forest fire monitoring: Enhancing patrol efficiency through risk assessment," *Fire*, vol. 7, no. 7, p. 254, 2024.

[25] C. Tymstra, R. Bryce, B. Wotton, S. Taylor, O. Armitage, *et al.*, "Development and structure of prometheus: the canadian wildland fire growth simulation model," *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Information Report NOR-X-417.(Edmonton, AB)*, 2010.

[26] S. Liu, L. He, and Z.-J. Max Shen, "On-time last-mile delivery: Order assignment with travel-time predictors," *Management Science*, vol. 67, no. 7, p. 4095–4119, July 2021.

[27] L. E. Ghaoui, M. Oks, and F. Oustry, "Worst-case value-at-risk and robust portfolio optimization: A conic programming approach," *Operations research*, vol. 51, no. 4, pp. 543–556, 2003.