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

DISASTER-AWARE INFORMATIVE PATH PLANNING IN EMERGENCY RESPONSE SCENARIOS

Xinya Cheng[1], Na Liu[1]

[1]School of Computer Science (National Pilot Software Engineering School),
Beijing University of Posts and Telecommunications, Beijing, China

chengxy@bupt.edu.cn, na.liu@bupt.edu.cn

ABSTRACT

In emergency response scenarios, rapid acquisition of critical disaster information supports effective decision-making. Traditional geometric coverage-based path planning often struggles to balance efficiency and information value. To address this, we propose a **Disaster-Aware Informative Path Planning (DAIPP)** method, which integrates a Siamese UNet-based building damage recognition model and formulates a novel information value function that considers recognition results, model uncertainty, and flight cost. We design an improved Frontier-based path planning algorithm, named **the Selective Frontier Algorithm (SFA)**, which enhances the selection of candidate points to achieve the prioritized exploration of critical regions. To validate its effectiveness, the proposed method is compared with coverage path planning, random planning, and Monte Carlo tree search (MCTS). Experiments on the xView2 dataset demonstrate that the proposed method outperforms baselines in terms of information coverage, semantic target hit rate, and weighted information coverage, providing strong support for efficient disaster perception in emergency response.

***Index Terms*—** Aerial Robotics, Informative Path Planning, Learning and Adaptive Systems

1. INTRODUCTION

The intensification of climate change and rapid urbanization has led to increasingly frequent natural disasters, posing serious threats to human lives and property. Due to the uncertainty of post-disaster environments and the time-critical nature of rescue operations, rapidly acquiring accurate high-value disaster area information has become essential. High-value disaster information refers to critical data such as damaged infrastructure, priority rescue areas, and regions with a high likelihood of secondary disasters. Traditional ground-based data collection methods are often constrained by terrain accessibility and communication limitations, resulting in low efficiency and limited coverage. In contrast, due to their high mobility and rapid deployment, unmanned aerial vehicles (UAVs) are regarded as an indispensable option to support emergency response [1, 2], including tasks such as generating real-time maps [3] and assisting communications [4].

When UAVs are deployed for emergency response operations, practical constraints, such as limited operational time and complex environmental conditions, often preclude complete sensing of the entire area. Consequently, it is critical to develop a resource-efficient path planning strategy that maximizes the acquisition of high-value information from critical zones. Currently, most disaster sensing methods still passively perform data acquisition, such as employing coverage path planning (CPP) [5, 6, 7]. Although effective for static scenarios, the CPP does not capture the information value of regions (e.g., areas prioritized for rescue, regions with a high likelihood of disaster occurrence) and lacks adaptability to dynamic disaster conditions. Informative path planning (IPP), which aims to maximize information gain, provides an effective solution. However, most IPP methods [8] highly depend on the accuracy of the prior model. If the real environment diverges from the model's predictions, which is very common in reality, the robot may end up collecting data along a suboptimal or even entirely ineffective path, failing to adapt to newly acquired information.

With the increasing complexity of tasks, researchers have proposed adaptive informative path planning (AIPP) [9], which is capable of dynamically adjusting paths based on real-time sensing data to actively and efficiently acquire high-value information in complex environments. For example, some studies focus on target navigation in unknown environments [10], while others address environmental monitoring [11, 12], image collection for model training [13, 14], and environmental information gathering [15, 16]. However, existing AIPP methods are not fully aligned with the requirements of emergency disaster response, as they rarely incorporate disaster-specific factors. The problem addressed in this work is to design a UAV path planning strategy that prioritizes the exploration of disaster-critical regions, thereby

This work was supported in part by the National Natural Science Foundation of China under Grant 62301038, and in part by the Fundamental Research Funds for the Central Universities under Grant No. 2025KYQD15, (*Corresponding author: N. Liu*).

enabling the rapid and efficient acquisition of actionable information to support time-sensitive rescue decision-making.

To address the core challenge of emergency response in acquiring high-value information with limited resources, we propose a disaster-aware informative path planning (DAIPP) method. As illustrated in Fig. 1, the framework integrates a Siamese UNet-based building damage recognition model with uncertainty evaluation to construct multi-layer information maps, and incorporates a selective frontier-based exploration strategy with a task-oriented information value function to guide the UAV in efficient and adaptive path planning. The contributions of this work are summarized as follows:

- We propose a disaster-aware informative path planning (DAIPP) framework to explore the high-value information in emergency response scenarios, establishing a perception–mapping–decision loop to prioritize disaster-critical regions and optimize UAV path planning;
- We design a dynamic multi-layer information map that combines building damage classification with uncertainty modeling for continuous decision-making;
- We formulate a task-oriented information value function integrating damage level, uncertainty, redundancy, and flight cost, coupled with the selective-frontier algorithm (SFA) for adaptive path planning.

2. RELATED WORK

2.1. Informative Path Planning

In traditional path planning tasks, the optimization objectives are typically shortest distance, minimal energy consumption, or maximum coverage. However, in information-gathering missions, particularly in post-disaster emergency response scenarios, the goal is no longer merely to cover all regions, but rather to acquire as much valuable information as possible within a limited time.

The core idea of IPP is that, at each decision point, the agent must select the path that maximizes information gain while accounting for traversal cost. With advances in robotics and the growing demand for adaptability, adaptive informative path planning (AIPP) has been developed on top of IPP. As described in [12], the general paradigm is formulated as follows:

$$\psi^* = \underset{\psi \in \Psi}{\operatorname{argmax}} I(\psi), \text{ s.t.}, C(\psi) \leq B, \tag{1}$$

where the observation position of the agent is defined as $p_i \in \mathbb{R}^3$, and Ψ denotes the set of all feasible positions. The objective is to maximize the information value $I : \Psi \rightarrow \mathbb{R}_{\geq 0}$ along the path $\psi = (p_1, \cdots, p_P) \in \Psi$. The total budget B can represent flight time, energy consumption, or other task constraints, while the cost function $C : \Psi \rightarrow \mathbb{R}_{\geq 0}$ measures the traversal cost of path ψ, i.e.,

$$C(\psi) = \sum_{i=1}^{P-1} c(p_i, p_{i+1}). \tag{2}$$

The choice of information metric is essential to AIPP. Existing studies have designed information value functions from different perspectives, including semantic class uncertainty with prediction probability [10], covariance differences from Kalman filtering combined with deep reinforcement learning [11], uncertainty reduction with multi-agent reinforcement learning [12], semantic segmentation uncertainty with data novelty [13], and Gaussian process-based metrics solved via approximate dynamic programming [15]. These approaches provide a solid theoretical foundation for intelligent information acquisition by UAVs in post-disaster response scenarios.

2.2. Uncertainty Estimation

Traditional path planning algorithms (e.g., $A^\star$, Dijkstra, and RRT) assume deterministic environments, yet real-world operations are inherently affected by uncertainties in perception, motion, and dynamic environments. Ignoring such uncertainties may lead to theoretically optimal but practically hazardous paths. Thus, uncertainty estimation becomes essential to derive robust and reliable paths that explicitly account for probabilistic risks. Estimating such uncertainty has been widely studied and plays a crucial role in path planning for information acquisition in unknown environments. Uncertainty estimation for deep learning models used in active exploration can be broadly categorized into two types: Bayesian neural networks [17] and ensembles [18].

Bayesian methods often employ the Monte Carlo dropout (MC dropout), performing multiple stochastic forward passes on a single model to approximate the posterior distribution. Ensemble methods, on the other hand, train multiple models and measure uncertainty based on their prediction variance. As reported in [19], Bayesian approaches generally incur higher computational costs due to repeated sampling, whereas ensemble methods offer a simpler and more efficient alternative.

In this work, motivated by the demand for the acquisition of high-value disaster information under constrained flight resources in emergency response scenarios, we employ a building damage classification model consisting of multiple heterogeneous networks. Under this design, the ensemble approach is more suitable for uncertainty estimation.

3. METHODOLOGY

As shown in Fig. 1, the overall framework of the proposed DAIPP consists of three main modules: disaster-aware in-

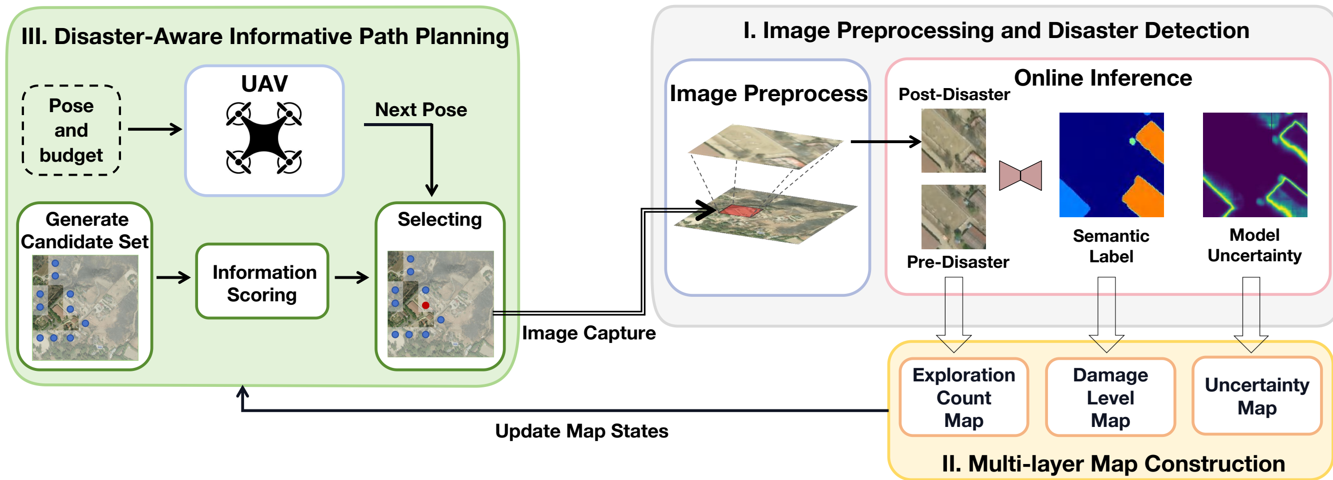

Fig. 1: Framework of the proposed DAIPP. The UAV first performs image preprocessing and online inference at the image processing and disaster detection module 3.1, generating semantic labels and model uncertainty estimates based on its current observations. The acquired information is then mapped into a three-layer representation 3.2. After the map is updated, it is utilized by the disaster-aware informative path planning module 3.3.

formative path planning for disaster perception, image preprocessing and detection of the disaster, and multilayer map construction. This method guides UAVs to prioritize exploration of areas with high uncertainty, obtain semantic and uncertainty information in real time, and dynamically update the global disaster map.

3.1. Image Preprocessing and Disaster Detection

This module simulates the acquisition of remote sensing imagery by UAVs during flight, maps them to the global map coordinate system, and then performs semantic reasoning on the remote sensing images collected within the UAV's current field of view to output the damage level of buildings and uncertainty estimates of the model.

In order to accurately project image information onto the map coordinate system, the system calculates the ground projection range using a geometric model based on the current flight altitude, camera field of view (FOV) parameters, and UAV attitude. It then converts this into a map grid index using ground resolution (GSD) to achieve precise alignment of the image with the map.

3.2. Multi-layer Map Construction

This module is used to record various types of information during the UAV's exploration process, providing spatial distribution references for the path planning module. We design three types of maps to optimize the UAV's exploration efficiency and accuracy.

We perform discretization of the task area into a two-dimensional regular grid $G \to \mathbb{R}^{W\times L}$, and construct the following three types of information layers for each grid: Dam-

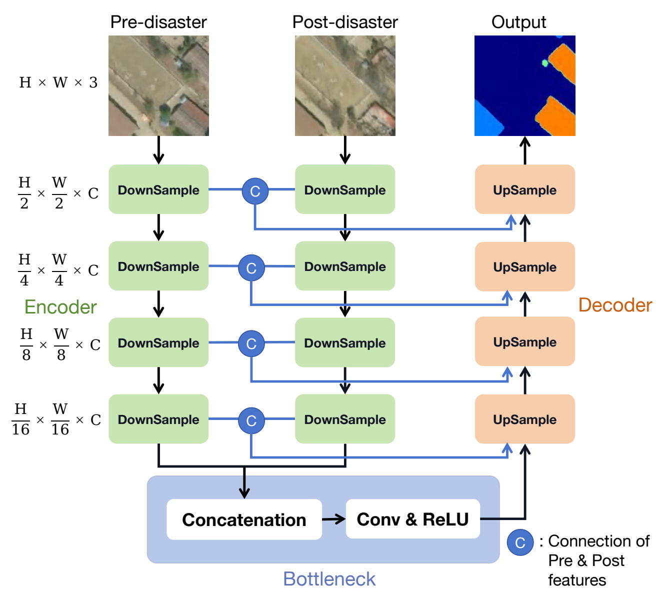

Fig. 2: Architecture of the image inference model. The Siamese UNet structure is used to locate buildings and classify damage levels, allowing the use of a variety of mainstream image classification networks as encoders, such as ResNet, DenseNet, DPN, etc.

age level map $\mathcal{G}_S :\to \{0, 1, \cdots, K-1\}^{W\times L}$, uncertainty map $\mathcal{G}_U :\to [0, 1]^{W\times L}$, and exploration count map $\mathcal{G}_E :\to N^{W\times L}$. When the UAV captures an image $\mathcal{Z}_t$ at position $p_t \in \mathbb{R}^3$, the online inference module outputs the building damage level label map $\mathcal{Y}_t$ and the corresponding model uncertainty score map $\mathcal{U}_t$. These images are mapped to the world coordinate system via a projection transformation, and then mapped to the map grid $\mathcal{G}^{m,n}$, updating the corresponding map layer at that location.

Damage Level Map: Using a grid-based approach, the damage level classification results for each area are recorded to visually reflect the disaster risks present in that area. Based on this, it is assumed that disaster risks are positively correlated with damage levels, meaning that areas with more severe damage have higher "exploration value." The update method involves directly writing the category labels obtained through online inference for the corresponding image areas into $\mathcal{G}_S$, overwriting previous records to ensure that the map always contains the latest damage assessment information.

Uncertainty Map: The uncertainty map $\mathcal{G}_U$ records the model's prediction uncertainty for each grid cell. Regions with higher uncertainty indicate insufficient model confidence and are therefore prioritized for further exploration. To estimate the uncertainty map $\mathcal{G}_U$, during each perception, K sub-models with different network structures and initialization parameters independently predict the images captured by the UAV. The prediction results $P_k(x,y)$ of these models at each grid (x,y) are treated as a set of samples, and their prediction standard deviation is used as the uncertainty score for that position, i.e.,

$$U(x,y) = \sqrt{\frac{1}{K}\sum_{k=1}^{K}\left(P_k(x,y) - \bar{P}(x,y)\right)^2}, \tag{3}$$

where $P_k(x,y)$ denotes the predicted probability of the damage class produced by the k-th model, and

$$\bar{P}(x,y) = \frac{1}{K}\sum_{k=1}^{K} P_k(x,y)$$

Exploration Count Map: The exploration count map $\mathcal{G}_E$ counts the times each grid cell has been explored by UAVs, serving as an indicator for assessing exploration redundancy. This map is used to penalize redundant exploration of the same location and encourage the UAV to focus on unexplored regions. Each time an area (x,y) is observed by a UAV, the exploration frequency is increased by one, updated as

$$\mu_{\mathrm{E}}^{t}(x,y) = \mu_{\mathrm{E}}^{t-1}(x,y) + 1. \tag{5}$$

3.3. Disaster-Aware Informative Path Planning

This module serves as the core component of the proposed framework. Its task is to generate a set of candidate observation locations based on the current multi-layer information map, calculate the scores of all candidate points based on information value modeling, and select the next exploration target position with the maximum information value.

3.3.1. Algorithm Pipeline

In the actual implementation process, the multi-layer map information is first updated based on the current observation position of the UAV. Then, a combination of candidate points is obtained using a candidate point screening strategy. All candidate points are traversed, and their information value is calculated sequentially to obtain a score for each candidate point. The point with the highest score is recorded as the current optimal target position until the area is fully explored or the set criteria are met. If all candidate points have a score of 0 (i.e., the entire area has been explored, or the model has not predicted any damaged areas), a point is randomly selected from the candidate points for flight to ensure the exploration task does not stall. To further ensure flight safety, the planner adds a secondary verification step before final confirmation of the target point.

3.3.2. Information Gain Modeling

We propose an information value function that integrates building damage severity information, model uncertainty scores, exploration redundancy penalties, and path costs to evaluate the "value" of each candidate observation point in path planning, thereby guiding UAVs to perform efficient information collection.

The information value function $I(x,y)$ is defined as,

$$\begin{aligned} I(x,y) = {} & \lambda_1 \cdot D(x,y) + \lambda_2 \cdot U(x,y) \\ & - \lambda_3 \cdot E(x,y) - \lambda_4 \cdot C(x,y) \end{aligned} \tag{6}$$

where $D(x,y)$ represents the damage level label of buildings in $\mathcal{G}_S$; $U(x,y)$ comes from the uncertainty map $\mathcal{G}_U$, indicating the model's prediction confidence for the current UAV observation area; $E(x,y)$ comes from the exploration count map $\mathcal{G}_E$, indicating the number of times the area has been observed. Increasing the exploration count as a penalty term helps avoid redundant exploration, while setting an additional base reward for unobserved areas encourages the UAV to prioritize exploring unobserved regions. $C(x,y)$ represents the path flight cost from the current position, encouraging the planning module to prioritize points with lower costs; λ_i are the weighting coefficients for each metric, used to adjust the importance of different factors in information assessment.

3.3.3. Candidate Point Selection Strategy

The Frontier algorithm [20] is a classical autonomous exploration method commonly used for robots or UAVs in unknown environments. Its core idea is to generate candidate target points at the interface between known and unknown regions (i.e., the frontier areas) to guide the agent continuously toward unexplored spaces. Compared with traditional global planning methods, this strategy offers strong real-time performance and scalability. However, the traditional Frontier algorithm has notable limitations: its path planning relies solely on grid occupancy information and selects the nearest reachable frontier point as the target among all detected frontiers. This approach lacks task-driven behavior, often results in ineffective or redundant exploration, and cannot meet the high-value information acquisition requirements of emergency rescue tasks. Some methods have improved this limitation. For example, [21] integrates object detection with Kalman filtering, constructing a semantically enhanced Frontier structure to introduce target location estimates into Frontier data, guiding UAVs to prioritize areas with potential targets.

To achieve path selection guided by information acquisition efficiency, this work builds upon the traditional Frontier candidate point generation by incorporating multi-layer map information and an information value function $I(x, y)$ to quantitatively evaluate potential information gain in frontier areas.

In this work, Frontier points are defined as contour points at the edges of the currently known map area. Each grid cell $\mu_E^{x,y}$ in the exploration map $\mathcal{G}_E :\to \mathbb{N}^{W\times L}$ is categorized as follows:

$$\text{Label}_{\mu_E(x,y)} = \begin{cases} \text{Known}, & \mu_E(x,y) > 0 \\ \text{Unknown}, & \mu_E(x,y) = 0 \end{cases} \tag{7}$$

Each observation position (x, y) has four neighboring cells (up, down, left, right) denoted as $\mathcal{N}(x, y)$. The candidate point set $\mathcal{F}$ is defined as:

$$\begin{aligned} \mathcal{F} = \{(x,y) \in \text{Known} \mid \exists (x',y') \in \mathcal{N}(x,y), \\ (x',y') \in \text{Unknown}\} \end{aligned} \tag{8}$$

Candidate points are sampled from all contour points in Known according to a step size δ to obtain the candidate set:

$$\mathcal{F}_{\text{candidate}} = \{(x_k, y_k) \in \mathcal{F} \mid k \equiv 0 \pmod{\delta}\} \tag{9}$$

To prevent the UAV from flying outside map boundaries or revisiting areas, constraints are added during candidate point selection: the candidate point must be within the map $\mathcal{G}$, its distance $d_p(x, y)$ to the current UAV position must exceed a minimum threshold $d_{\min}$, and its flight cost $C_p(x, y)$ must be less than the total flight budget B. The resulting valid candidate set is:

$$\begin{aligned} \mathcal{F}_{\text{valid}} = \{(x,y) \in \mathcal{F}_{\text{candidate}} \mid\ & (x,y) \in \mathcal{G}, \\ & d_p(x,y) > d_{\min}, \\ & C_p(x,y) < B\} \end{aligned} \tag{10}$$

4. EXPERIMENT

To validate the effectiveness of the UAV information collection path planning method proposed in this work, we conducted simulation experiments based on the remote sensing public dataset xView2[1] to quantitatively evaluate and visually analyze the path planning performance and information acquisition efficiency of the system. This section mainly includes experimental settings, evaluation index descriptions, and path planning performance comparisons.

(a) Pre-disaster (b) Post-disaster (c) Overlay

Fig. 3: This example shows images of the SoCal region before and after the fire. The sample consists of two images taken from the same location before the fire 3a and after the fire 3b. In 3c, red represents damaged buildings, orange represents severely damaged buildings, blue indicates buildings with minor damage (not included in the example), and green indicates undamaged buildings.

4.1. Dataset Description

The xView2 dataset provides a large number of pre- and post-disaster remote sensing image pairs, accompanied by building contour and damage level labels, making it suitable for post-disaster assessment and path simulation tasks. Considering the continuity and coverage requirements of path planning systems, we select city-level scenes from the dataset as test areas to ensure that the image content has characteristics such as uneven building density and varying degrees of damage.

The dataset example is shown in Fig. 3. Each sample in the dataset consists of a pair of pre-disaster and post-disaster images, forming a dual-temporal image pair. The image resolution is 0.3 meters/pixel, with high image quality and clear object boundaries, making it suitable for downstream tasks such as building identification and damage assessment. Each image is accompanied by vector boundaries of buildings and damage severity labels for each building. The severity levels are as follows: 0 represents non-building areas, 1 represents non-damaged building areas, 2 represents minor damaged building areas, 3 represents major damaged building areas, and 4 represents completely destroyed building areas.

4.2. Evaluation Metrics

The existing path planning algorithm evaluation metrics for UAV emergency information collection tasks have certain limitations. Therefore, we design a two-dimensional evaluation framework that includes value-aware coverage and semantic hierarchical response.

Weighted Information Coverage Rate: This metric is used to measure the coverage of all information-rich areas (i.e., building areas) in post-disaster scenarios by UAVs with

[1]https://xview2.org/

limited exploration attempts or flight budgets. By assigning different weights to buildings with different damage levels, it reflects whether path planning effectively guides UAVs to focus on information-rich areas, which is the core evaluation dimension of this work. When $\omega_c(x, y) \equiv 1$, it degenerates into a traditional coverage metric.

Let the total information value area in the entire map be I_{total}, the weighted value area actually covered by the UAV during exploration be $I_{obtained}$, the pixel point within the building area be (x, y), its semantic damage level be $c(x, y)$, and the semantic weight $\omega_{c(x,y)}$ corresponding to each level. the number of UAV explorations is $H(x, y)$, then the information coverage rate is defined as:

$$\begin{aligned} WICR &= \frac{I_{obtained}}{I_{total}} \times 100\% \\ &= \frac{\sum_{(x,y)\in\mathcal{B}} \omega_{c(x,y)} \cdot \min(1, H(x, y))}{\sum_{(x,y)\in\mathcal{B}} \omega_{c(x,y)}} \end{aligned} \tag{11}$$

Semantic Target Hit Rate: This indicator is used to measure whether the UAV has successfully explored a specific semantic target area. Combined with the semantic damage level of buildings, it assesses whether path planning can prioritize high-priority areas. The calculation formula is as follows:

$$STHR^c = \frac{B^c_{obtained}}{B^c_{total}} \times 100\%, \quad c \in \{1, 2, 3, 4\} \tag{12}$$

Where B^c_{total} denotes the total number of pixels for buildings with semantic category c, and $B^c_{obtained}$ denotes the number of pixels successfully explored in buildings with semantic category c.

4.3. Performance Analysis of Path Planning Strategies

We use the example in Fig. 3 for experiments, comparing five path - planning algorithms with distinct mechanisms: the proposed method (AIPP Frontier), driven by multi - layer maps and information value functions to integrate semantic and value - based guidance; the MCTS algorithm, which leverages Monte Carlo tree search, treating information value functions as simulated rewards to select paths by simulating multiple routes and picking the highest - expected - reward direction (with step and node expansion constraints); the Traditional Frontier, a simpler approach relying solely on boundary point distance for selection, lacking semantic or value - driven logic; Coverage Path Planning (CPP), which enforces full environmental coverage via geometric patterns like zigzags; and the Random algorithm, which haphazardly chooses feasible directions at each step with no deliberate decision framework.

The information value of a task is directly proportional to the damage level of a building. In UAV path planning, high-priority collection of areas with high information value should be reflected in prioritizing the collection of buildings with higher damage levels. We set the flight cost to 5000 units. As shown in Fig. 4a, the path planned by the AIPP Frontier algorithm proposed in this work primarily surrounds the building areas in the scene, while the areas in the lower left and lower right of the figure that do not contain buildings are not selected for observation, thereby avoiding the consumption of flight resources. MCTS, as an intelligent search algorithm 4b, although it possesses certain global search capabilities, its performance is inferior to that of AIPP Frontier. As shown in Fig. 4c, the traditional Frontier algorithm does not prioritize buildings with high damage levels, failing to achieve intelligent information collection. The CPP algorithm, as shown in Fig. 4d, simply performs information collection in the scene using geometric search methods. The Random algorithm, as shown in Fig. 4e, also lacks the ability to identify high-value information areas.

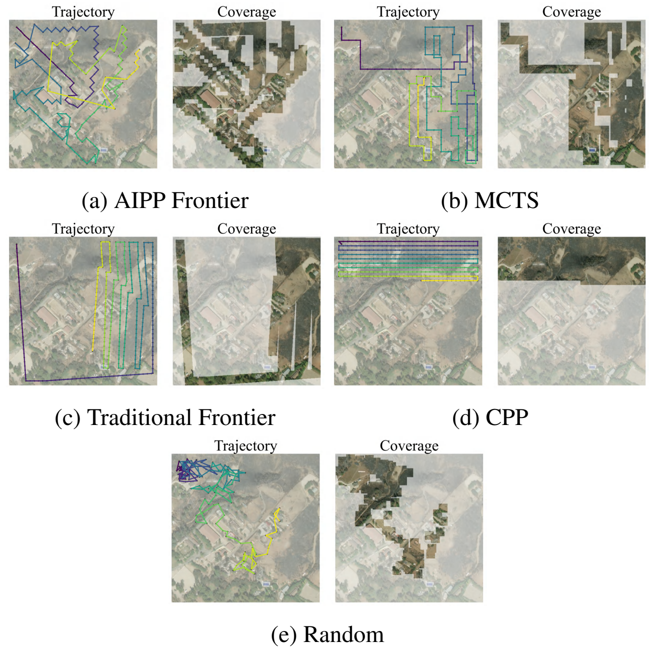

Fig. 4: Path planning results for five algorithms under a flight budget of 5,000 units. Opaque areas represent regions that the UAV has not explored, while completely clear areas represent regions that the UAV has explored more than twice (in this work, the UAV is assumed to sense in the center of the region, so there will be overlap in the field of view).

In order to further analyze the performance of the five algorithms, we use the traditional information collection evaluation index of information coverage, combined with the two indicators designed in 4.2, to comprehensively analyze the five algorithms.

Based on the experimental results shown in Fig. 5: Fig. 5a, which illustrates how information coverage varies with flight budget, the five algorithms exhibit significant differences: The AIPP Frontier algorithm designed in this work performs optimally, achieving a final information coverage rate of 92%, and demonstrates the highest efficiency under

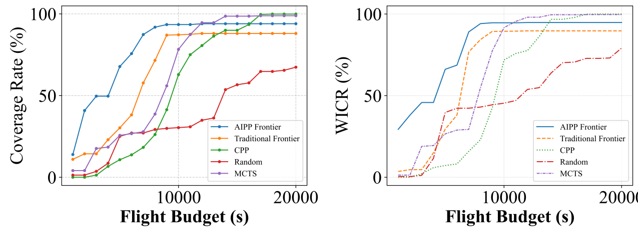

(a) Coverage Rate (b) WICR

Fig. 5: Coverage and information coverage results for the five algorithms. The steeper the curve, the better the performance. The AIPP Frontier algorithm (blue) designed in this work performs best, enabling priority sensing of high-value areas with limited resources.

limited flight budgets (reaching 80% coverage at an 8,000 budget), enabling rapid collection of most information; The traditional Frontier algorithm follows next, with lower initial efficiency compared to the proposed method and ultimately achieving only 85% information coverage, reflecting the inherent drawback of traditional methods prone to local optima. The remaining three algorithms lag significantly behind the proposed method.

As shown in Fig. 5b, the AIPP Frontier algorithm proposed in this work significantly outperforms the other four strategies in terms of weighted information coverage. Compared to the traditional Frontier method, the method proposed in this work exhibits a faster growth curve when the flight budget is less than 5,000, indicating that it can efficiently capture high-value areas in the early stages. As the budget increases, the algorithm's weighted coverage continues to rise steadily and reaches saturation at approximately 12,500 budget, with the optimal value approaching 95%. This result indicates that in resource-constrained real-world tasks, the proposed Frontier algorithm can achieve maximum information gain at a limited exploration cost.

From the results shown in Fig. 6, the proposed AIPP Frontier algorithm shows significantly better performance in terms of the semantic hit rate at different levels of building damage compared to the other four strategies. Compared to the traditional Frontier method, in high-value areas such as severely damaged and destroyed buildings, the AIPP Frontier algorithm exhibits a steeper growth curve even before the flight budget reaches 7500. This indicates that, at an early stage, it can more efficiently focus on high-value, heavily damaged areas and prioritize the acquisition of critical information. As the flight budget increases, the semantic hit rate of the AIPP Frontier algorithm in regions of serious damage continues to increase steadily. In severely damaged regions, the hit rate rapidly reaches 100% when the budget approaches 7500; in the destroyed regions, the hit rate also climbs quickly to nearly 100% at around the same budget and then remains stable.

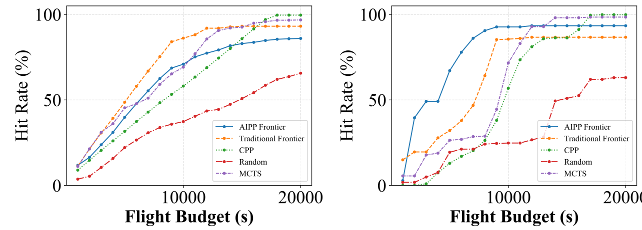

(a) No Building (b) No Damaged

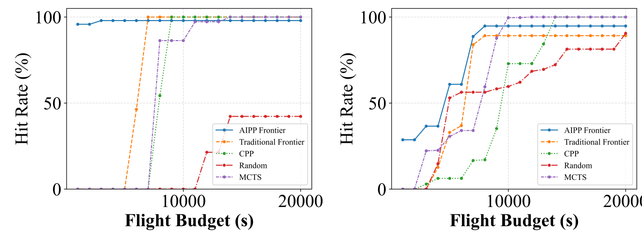

(c) Major Damaged (d) Destroyed

Fig. 6: The hit rates of the five algorithms in different categories of areas (the data used does not include Minor Damaged areas). The steeper the curve, the better the hit rate of the algorithm in that category of area.

5. CONCLUSIONS

This work investigates UAV path planning for intelligent information collection in emergency rescue scenarios. We propose a perception-driven system within the Informative Path Planning (IPP) framework, which integrates semantic understanding, uncertainty modeling, multi-layer mapping, and frontier-based strategies to prioritize high-value regions. Experimental results on the xView2 dataset demonstrate improved information coverage and reduced redundancy, with visualizations highlighting effective focus on disaster-critical areas. Nevertheless, challenges remain in generalization, mapping errors, local optimality, and manual design of information value functions. Moreover, the current evaluation is conducted in simulation based on public datasets, and systematic validation in real-world disaster environments has not yet been performed. Future work will explore domain adaptation, visual SLAM, global optimization, and policy learning, alongside real-world UAV deployment for practical validation.

6. REFERENCES

[1] Milan Erdelj and Enrico Natalizio, "Uav-assisted disaster management: Applications and open issues," in *2016 international conference on computing, networking and communications (ICNC)*. IEEE, 2016, pp. 1–5.

[2] Hafsa Ahmed, Mahmoud Bakr, Manar Abu Talib, Sohail Abbas, and Qassim Nasir, "Unmanned aerial vehicles (uavs) and artificial intelligence (ai) in fire related

disaster recovery: analytical survey study," in *2022 International Conference on Business Analytics for Technology and Security (ICBATS)*. IEEE, 2022, pp. 1–6.

[3] Helen Oleynikova, Christian Lanegger, Zachary Taylor, Michael Pantic, Alexander Millane, Roland Siegwart, and Juan Nieto, "An open-source system for vision-based micro-aerial vehicle mapping, planning, and flight in cluttered environments," *Journal of Field Robotics*, vol. 37, no. 4, pp. 642–666, 2020.

[4] Francesco Malandrino, Carla-Fabiana Chiasserini, Claudio Casetti, Luca Chiaraviglio, and Andrea Senacheribbe, "Planning uav activities for efficient user coverage in disaster areas," *Ad Hoc Networks*, vol. 89, pp. 177–185, 2019.

[5] Tauã M Cabreira, Lisane B Brisolara, and Ferreira Jr Paulo R, "Survey on coverage path planning with unmanned aerial vehicles," *Drones*, vol. 3, no. 1, pp. 4, 2019.

[6] Haiyang Lu, Yuhao Yang, Rentuo Tao, and Yawei Chen, "Coverage path planning for sar-uav in search area coverage tasks based on deep reinforcement learning," in *2022 IEEE International Conference on Unmanned Systems (ICUS)*. IEEE, 2022, pp. 248–253.

[7] Hongyu Song, Jincheng Yu, Jiantao Qiu, Zhixiao Sun, Kuijun Lang, Qing Luo, Yuan Shen, and Yu Wang, "Multi-uav disaster environment coverage planning with limited-endurance," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10760–10766.

[8] Halit Ergezer and Kemal Leblebicioglu, "Path planning for uavs for maximum information collection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 502–520, 2013.

[9] Marija Popović, Joshua Ott, Julius Rückin, and Mykel J Kochenderfer, "Learning-based methods for adaptive informative path planning," *Robotics and Autonomous Systems*, vol. 179, pp. 104727, 2024.

[10] Georgios Georgakis, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, and Kostas Daniilidis, "Learning to map for active semantic goal navigation," *arXiv preprint arXiv:2106.15648*, 2021.

[11] Julius Rückin, Liren Jin, and Marija Popović, "Adaptive informative path planning using deep reinforcement learning for uav-based active sensing," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4473–4479.

[12] Jonas Westheider, Julius Rückin, and Marija Popović, "Multi-uav adaptive path planning using deep reinforcement learning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 649–656.

[13] Julius Rückin, Federico Magistri, Cyrill Stachniss, and Marija Popović, "An informative path planning framework for active learning in uav-based semantic mapping," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4279–4296, 2023.

[14] Julius Rückin, Liren Jin, Federico Magistri, Cyrill Stachniss, and Marija Popović, "Informative path planning for active learning in aerial semantic mapping," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11932–11939.

[15] Joshua Ott, Edward Balaban, and Mykel J Kochenderfer, "Trajectory optimization for adaptive informative path planning with multimodal sensing," in *2024 10th International Conference on Control, Decision and Information Technologies (CoDIT)*. IEEE, 2024, pp. 1–7.

[16] Mariliza Tzes, Nikolaos Bousias, Evangelos Chatzipantazis, and George J Pappas, "Graph neural networks for multi-robot active information acquisition," *arXiv preprint arXiv:2209.12091*, 2022.

[17] Yarin Gal and Zoubin Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.

[18] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.

[19] Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al., "A survey of uncertainty in deep neural networks," *Artificial Intelligence Review*, vol. 56, no. Suppl 1, pp. 1513–1589, 2023.

[20] Brian Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97.'Towards New Computational Principles for Robotics and Automation'*. IEEE, 1997, pp. 146–151.

[21] Hai Lin, Xinsong Yang, Guanghui Wen, and Wei Xing Zheng, "Fast uav object-searching in large-scale and complex environments," *IEEE Transactions on Cybernetics*, 2025.