

Article

Ranking Inversion in Risk-Parameterised UAV Path Planning for Wildfire Emergency Response

Konstantinos Zervakis and Ilias Panagiotopoulos *

Department of Military Sciences, Hellenic Army Academy, Evelpidon Avenue,
166 73 Vari, Greece; zervakisai@gmail.com

* Correspondence: ipanagiotopoulos@sse.gr

Abstract

Wildfires generate rapidly evolving hazard landscapes that disrupt ground-based logistics and render conventional disaster-response operations ineffective, motivating the use of unmanned aerial vehicles (UAVs) in civil-protection missions such as medical resupply, casualty search-and-rescue, and perimeter surveillance. Existing evaluations, however, share a common limitation: they assess performance using navigation-centric metrics—primarily success rate—without accounting for the temporal value of the mission objective. This paper characterises, within each planner family, how a single risk coefficient q governs the trade-off between navigation success and time-decaying mission value: holding each family's algorithm and replan trigger fixed and sweeping only q isolates its effect, so that the risk setting maximising a family's navigation success need not maximise its mission value. To study this, FLARE is introduced, a deterministic benchmark for fire-landscape adaptive risk evaluation, inspired by recent Greek wildfire events; all hazard dynamics (fire spread, structural collapse, moving obstacles, dynamic no-fly zones) are modelled as benchmark abstractions rather than incident-specific reconstructions. FLARE evaluates eight planners spanning six risk-handling paradigm families, isolating the risk coefficient by sweeping q within each family (algorithm fixed) with A* as the risk-blind static reference, and quantifies mission impact—medication efficacy, casualty survival, and data freshness—through a strictly time-decreasing mission-score function. The mission-value leverage of q is strongly family-dependent: decisive for the incremental soft-cost-inflation family—whose success-optimal q collapses its mission value ($0.91 \rightarrow 0.41$)—strong for the worst-case (CVaR) family, moderate for the sampling family, and negligible for the reactive, hard-threshold and frequent-replan families. In some families the success-optimal and mission-optimal q diverge—a within-family ranking inversion—while in others they coincide; because the algorithm is fixed across each sweep, the divergence is attributable to q alone. Success of navigation is therefore a necessary but insufficient proxy for mission effectiveness, and the risk coefficient is a meaningful per-family parameter for mission-aware configuration.

Academic Editor: Heng Shi

Received: 18 June 2026

Revised: 21 July 2026

Accepted: 31 July 2026

Published: 7 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

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

Keywords: UAV path planning; wildfire emergency response; ranking inversion; risk-aware planning; mission-impact evaluation; deterministic benchmarking; reproducibility

1. Introduction

Natural disasters such as wildfires, earthquakes, and floods generate rapidly evolving hazard landscapes that severely disrupt ground-based logistics. In such events, the timely delivery of medical supplies, situational awareness, and casualty location constitutes the primary operational objective for emergency responders [1,2]. Traditional relief operations rely on manned helicopters and ground convoys; however, both modalities suffer from road network damage, airspace congestion, and limited scalability [3]. Unmanned aerial vehicles (UAVs) have emerged as a promising alternative, offering low-cost, rapid-deployment platforms capable of operating in degraded environments where manned access is infeasible [4,5].

Over the past decade, extensive research has addressed the UAV path planning problem across a range of algorithmic families: graph-search methods such as A* [6] and D* Lite [7], potential-field approaches [8], sampling-based planners, meta-heuristic optimisers, and, more recently, deep reinforcement learning agents [5,9,10]. Comprehensive surveys confirm that path planning remains one of the most actively studied problems in UAV autonomy, with over 3500 publications indexed between 2000 and 2024 [11,12]. Risk-aware extensions that modulate edge costs by proximity to hazards have further improved planner robustness in cluttered environments [13–15].

Despite these advances, most existing evaluations share a common limitation: they assess planner performance using navigation-centric metrics—path length, success rate, or computation time—without accounting for the temporal value of the mission objective [16]. In disaster response, however, mission effectiveness degrades monotonically with elapsed time: insulin delivered after 200 steps carries substantially greater clinical value than the same delivery at 800 steps [17,18]. Navigation success is therefore a necessary but insufficient proxy for mission effectiveness.

This distinction is particularly acute in wildfire events, where the hazard landscape expands continuously, and cellular-automaton (CA) fire models with wind-directional spread [19–23] have been widely adopted to capture this geometry.

In light of the above, the present study characterises, within each planner family, how the risk coefficient q mediates the trade-off between navigation success and time-decaying mission value. To investigate this effect, FLARE (Fire-Landscape Adaptive Risk Evaluation) is introduced, a deterministic benchmarking framework in which each family is parameterised by the risk coefficient applied to a shared hazard cost map and evaluated against a time-dependent mission-value function that jointly captures survival probability and time-induced degradation from a single deterministic replay. To the best of our knowledge, no prior work integrates a risk-parameterised planner configuration with a strictly time-decreasing mission-value function in a deterministic and reproducible wildfire benchmark. FLARE's specific contribution is to show, in this controlled setting, that within a single planner family the navigation-optimal risk coefficient can diverge from the mission-optimal one.

The primary contributions of this study are threefold.

(a) First, a deterministic benchmarking framework for wildfire emergency response is introduced for the rigorous evaluation of UAV path planners, considering coupled time-varying hazard dynamics (cellular-automaton fire spread, structural collapse, road-network interdictions), and situated against existing simulators through a four-dimensional capability matrix.

(b) Second, a risk-parameterised planner suite is presented in which a single scalar q makes risk sensitivity an explicit per-family parameter. The core instrument is a within-family sweep of q with each algorithm held fixed, so that any resulting change in performance is attributable to the risk coefficient alone under identical environmental conditions, directly isolating q from planner architecture.

(c) Third, the within-family ranking inversion is characterised: for each planner family, how its own optimal risk coefficient q shifts across metrics (navigation success, survival probability, mission effectiveness) is quantified, establishing that the inversion is a per-family regime—decisive in some families and absent in others.

The remainder of this paper is organised as follows: Section 2 reviews the relevant literature and positions FLARE among existing simulators. Section 3 introduces the benchmark design and the proposed approach. Section 4 presents the per-family risk-coefficient sweep, the resulting within-family ranking inversion, and its robustness. Section 5 interprets the main findings. Finally, Section 6 concludes the paper and suggests avenues for future research.

2. Background and Related Work

This section provides an overview of the fundamental concepts and prior research relevant to the present study with respect to UAV path planning in disaster responses, reproducible benchmarking, and dynamic environments.

The operational deployment of UAVs in disaster-response contexts has been the subject of several recent comprehensive reviews. Yücesoy et al. [2] survey operations-research applications of drone technology across the disaster-management cycle, complemented by reviews of search-and-rescue [24] and healthcare and emergency services [25].

Path planning remains a central enabling technology for these missions. Recent taxonomies classify planning algorithms into five principal categories—classical, heuristic, meta-heuristic, machine-learning, and hybrid methods [26], and ablation studies on hybrid architectures have been reported in [13,27–30]. Additionally, reproducibility standards for learning-based planners originate from Henderson et al. [12] and have been extended in the broader benchmarking literature [31–34].

With respect to the challenges associated with dynamic fire environments, CA models with wind-directional spread [19–21] provide the base dynamics on top of which Wang et al. [22] address multi-target firefighting task allocation in dynamic forest fire environments, while Liu et al. [23] and Zhao et al. [35] consider collaborative swarm strategies under resource constraints.

Within graph-search planning, A^* [6] yields cost-optimal paths under admissible heuristics on static graphs but requires a full re-expansion whenever the cost map changes, which is incompatible with the per-step hazard updates produced by wildfire dynamics. D^* Lite [7] addresses this limitation through incremental graph repair, reusing the search tree across replans and bounding worst-case re-expansion to the cells whose costs changed since the last query. Anytime variants such as ARA^* and the lifelong-planning family LPA^* further trade optimality for bounded responsiveness, but in all cases the planner remains agnostic to the temporal value of the mission objective—the metric of merit is path length or replan latency, not the value of the goal at the time of arrival.

Risk-aware extensions of these planners modulate edge costs by proximity to hazards [13–15,29] and have been applied to drone routing under obstacle uncertainty, GPS-degraded environments, and chemical-plume navigation. The common construction inflates traversal cost as a monotonically increasing function of a distance-decayed risk field. None of these formulations, however, couples the risk inflation to a strictly decreasing mission-value function: the implicit assumption is that completion is binary (success/failure), so longer detours that retain feasibility are preferred over shorter routes that admit residual hazard.

Moreover, surveys of UAV path planning [16] note that the dominant evaluation criterion remains a navigation-centric scalar—success rate, path length, or computation time—with mission-level objective value absent from the standard reporting protocol. We argue that this gap is the structural mechanism behind the within-family ranking

inversion: when a single planner is swept over its risk coefficient q , the binary success metric and a strictly decreasing mission score (which discounts delivery time) can be optimised at different values of q , so the risk setting that maximises a planner's navigation success need not maximise its mission value.

Following the above, although significant progress has been made in UAV path planning and reproducible benchmarking, existing evaluation frameworks predominantly rank planners by navigation-centric metrics and rarely couple a risk-parameterised planner to a strictly time-decreasing mission value within a single reproducible benchmark. These limitations motivate the need for further investigation. In this context, our work builds upon and extends prior efforts by addressing these challenges through the introduction and development of FLARE, thereby contributing to a more effective and comprehensive solution. By coupling a strictly decreasing mission score with expanding hazard dynamics, FLARE reveals performance trade-offs that success rate alone cannot capture.

Mission-oriented evaluation is also established in adjacent operational studies: drone-delivered defibrillator studies measure time-to-patient gains [36,37], the Rwandan national drone network is assessed by blood-product delivery time and wastage [38], and the orienteering literature formalises time-dependent prize collection [39,40]. These studies quantify mission value operationally, but each evaluates a fixed routing design in isolation; none couples risk-parameterised planner configurations to a strictly time-decaying mission value inside a deterministic, replayable hazard simulation, which is the gap FLARE addresses.

To situate FLARE among existing tools, Table 1 compares representative UAV, robotics, and wildfire simulation environments across the four capability dimensions that the ranking inversion requires: coupled hazard modelling, a strictly time-decreasing mission value, a unified risk coefficient applied identically across planners, and bit-identical deterministic replay. Most existing environments model only vehicle dynamics or a single hazard class [41–46] and treat task completion as binary, so within a single planner, its navigation-optimal and mission-optimal risk settings cannot diverge. The closest related evaluation, Kelly and Chaudhry [47], likewise observes that navigation cost and mission value need not agree; FLARE differs in combining coupled, interacting hazard dynamics, an explicit time-decaying mission value, a single risk coefficient shared across planner families, and seeded deterministic replay in one environment, which is what allows the ranking inversion to be reproduced and attributed.

Table 1. Comparison of FLARE with representative UAV, robotics, and wildfire simulation environments across the four capability dimensions underlying the ranking inversion, plus the resulting ranking-inversion column. ‘Partial’ denotes limited or implicit support.

	Coupled Hazards	Time-Decay Mission	Unified q	Deterministic Replay	Ranking Inversion
AirSim [41]	No	No	No	Partial	No
Flightmare [42]	No	No	No	Partial	No
gym-pybullet-drones [43]	No	No	No	Partial	No
FireCommander [44]	Partial	Partial	No	Partial	No
CREW-Wildfire [45]	Partial	Partial	No	Partial	No
FARSITE [46]	Partial	No	No	Partial	No
Kelly and Chaudhry [47]	Partial	Partial	Partial	No	Yes
FLARE (present work)	Yes	Yes	Yes	Yes	Yes

3. Materials and Methods

This section describes the FLARE framework, which is used to evaluate UAV risk-aware planning under dynamic hazard conditions. Figure 1 presents the overall architecture of the FLARE framework, which consists of five principal components:

- (i) an OpenStreetMap-derived grid environment;
- (ii) a coupled hazard dynamics module comprising wind-driven fire, structural collapse, and dynamic traffic obstacles with mutual interactions;
- (iii) a two-stage perception pipeline that fuses hazard layers into a continuous risk map $R(x)$ and inflates edge costs via the risk coefficient ρ ;
- (iv) a risk-parameterised planner suite evaluated per family: the risk coefficient ρ is swept within each family while its algorithm and replan trigger are held fixed, across eight planners spanning six paradigm families, with A* retained as the risk-blind static reference; and
- (v) a deterministic runner that orchestrates episode execution, mission scoring via strictly decreasing functions $M(t)$, and statistical evaluation. Reproducibility is enforced through a number of automated contracts. FLARE (version 1.0.0) is implemented in Python (version 3.14) and builds on NumPy (version 2.4), SciPy (version 1.17), NetworkX (version 3.6), Matplotlib (version 3.10), Gymnasium (version 1.2) [48] for the environment interface, and OSMnx (version 2.0) [49] for the extraction of the street networks.




Figure 1. FLARE benchmark main architecture. The wind-driven fire layer follows the cellular-automaton spread model of Alexandridis et al. [19].

Additionally, FLARE adopts three specific protocols: the interquartile mean (IQM) with bootstrapped confidence intervals from Agarwal et al. [31] for robust central-tendency estimation; Bonferroni-corrected Wilcoxon signed-rank tests as the non-parametric pairwise discriminator; and a paired-seed design in which the same RNG seed is replayed across all the planners so that seed-level variance cancels in the paired differences and only planner-level variance enters the test statistic.

3.1. Environment

FLARE operates on a grid $G = \{0, \dots, N\}^2$ extracted from OpenStreetMap tiles [48,49]. The action-space cardinality is $|A| = 5$, denoting the size of the discrete action space $A = \{\text{north, south, east, west, hold}\}$. The spatial connectivity of the grid remains 4-connected because hold does not change the cell. This 4-connected choice is a standard design

decision for grid-based planning benchmarks; it is adopted here because (i) it admits exact optimality for A* with Manhattan heuristic and (ii) it preserves comparability with prior grid planners.

Diagonal (8-connected) moves could be enabled in FLARE by changing a single configuration flag, at the cost of losing the exact-optimality guarantee and changing the cost-scale calibration. A strict timing contract governs each step: the planner observes the current blocking mask (frozen), selects an action, and only then does dynamics advance. If fire or debris reaches the agent's cell after the dynamics update, the episode terminates immediately.

3.2. Hazard Dynamics Module

Three coupled layers produce a hazard landscape module that worsens monotonically. This includes:

Wind-driven fire. Spread follows a cellular automaton on the Moore (8-cell) neighbourhood, modulated by wind direction and land use. Downwind cells ignite faster; upwind cells are suppressed. A buffer zone and dense smoke further restrict UAV movement. The wind parameter introduces directional asymmetry in the hazard geometry: isotropic fire expands symmetrically, whereas wind-driven fire elongates along the downwind axis, selectively blocking corridors while leaving upwind flanks traversable.

Structural collapse. Buildings exposed to fire for a threshold duration τ_c (the fire-exposure time before structural failure) collapse with probability p_d (the collapse likelihood upon exceedance of τ_c), scattering permanent debris. Unlike fire—which burns out—debris accumulates, progressively eliminating detour options.

Dynamic obstacles. Moving obstacles patrol the road and maritime network, blocking UAV movement within a Manhattan (L1, taxicab) buffer of configurable radius. For an airborne UAV, these model low-altitude corridor stand-off constraints—exclusion buffers over moving emergency vehicles and vessels under which low-flying UAVs must not operate. The interaction engine couples fire to traffic: fire proximity closes segments, creating compound blockages that neither layer produces in isolation.

3.3. Risk Cost Map

The perception pipeline maps raw hazards to risk-weighted edge costs in two stages. Let $R(x): G \rightarrow [0,1]$ denote a continuous risk map that fuses multiple hazard layers into a single spatial risk field at each grid cell $x \in G$, and let $w(x)$ denote the resulting weighted edge cost used by each planner. The two-stage formulation follows established risk-aware path planning practice [29], combining risk fusion and cost inflation.

Stage 1: Risk fusion. Five distance-decay layers (fire, smoke, traffic, debris, no-fly zones) are fused into a continuous risk map $R: G \rightarrow [0,1]$ via element-wise max:

$$R(x) = \max(R_f, R_s, R_t, R_b, R_{nfz}) \quad (1)$$

with the five layer-specific risk fields defined as the sub-equations below:

$$\text{Fire: } R_f(x) = [1 - d_f(x)/r_f]^+ \quad (1a)$$

$$\text{Smoke: } R_s(x) = s(x)/2 \quad (1b)$$

$$\text{Traffic: } R_t(x) = c_t \cdot [1 - d_t(x)/r_t]^+ \quad (1c)$$

$$\text{Debris: } R_b(x) = c_b \cdot [1 - d_b(x)/r_b]^+ \quad (1d)$$

$$\text{No-fly zone: } R_{nfz}(x) \in \{0, 0.5, 1\} \quad (1e)$$

for every cell $x \in G$. Here df , dt , db are Euclidean distance transforms from fire, traffic, and debris cells, respectively, with influence radii rf , rt , rb ; $s \in [0,1]$ is the smoke concentration; ct and $cb \in [0,1]$ are the per-layer severity coefficients; and $[\cdot]^+ = \max(\cdot, 0)$. Specific numeric values for rf , rt , rb , ct , and cb are reported per scenario in Section 4.2. The per-layer weights ct and cb are chosen so that, at distance zero, a traffic obstacle is slightly less restrictive than a fire cell and a debris cell is slightly more restrictive than a traffic cell; together these scaling constants (ct , cb) are calibrated heuristics reflecting the relative severity of each hazard layer. We use max rather than summation so that a single severe hazard dominates regardless of benign neighbours.

Stage 2: Cost inflation. Each family inflates its edge cost using the risk coefficient $q \in \mathbb{R} \geq 0$, through the following weight function

$$w(x) = 1 + \rho \cdot R(x) \quad (2)$$

The risk coefficient q —rather than the algorithmic family—determines the effective cost landscape: an identical hazard state produces distinct cost surfaces as q varies, from uniform cost ($q = 0$, risk-blind) to near-impassable barriers (high q , risk-averse). High q amplifies perceived danger, causing wider detours, whereas low q allows routes through fire fringes. Because q is a common risk-aversion coefficient—the cost-inflation law, Equation (2), is applied in identical form by the graph-search families, and the potential-field family applies an analogous risk-scaled repulsive gain—it can be swept within any family while the algorithm is held fixed; this within-family sweep is described in detail in Section 4.3, and no per-family value is fixed a priori.

3.4. Planner Suite

The planner suite is organised into six risk-handling paradigm families and studied by sweeping the risk coefficient q within each family while its algorithm and replan trigger are held fixed. Eight planners span the six families; each family is governed by the single scalar q that modulates its sensitivity to the shared hazard risk cost map $R(x)$, which makes risk preference explicit and uniformly defined across families under identical environmental conditions.

Crucially, no per-family risk coefficient is fixed a priori: q is the object of study, swept within each family (holding the algorithm and replan trigger fixed) so that every performance statement in Section 4 is made per planner family rather than as a cross-planner ranking at a fixed q . A^* [6] is retained only as the risk-blind static reference—it ignores the risk map and is excluded from the sweep. The six paradigm families are:

- **Soft cost-inflation:** Periodic (fixed-interval replan), Aggressive (obstacle-mask-change replan), and Incremental A^* (D^* Lite-inspired incremental repair [7]).
- **Hard threshold:** a chance-constrained planner with a hard risk bound [50].
- **Worst-case/CVaR:** a conservative planner that inflates cost by a tail-risk measure [51,52].
- **Sampling:** a risk-aware RRT* planner [53].
- **Exponential/entropic:** a risk-sensitive planner with multiplicative cost $\exp(q \cdot R)$ [54,55].
- **Potential field:** APF [8], applying a risk-scaled repulsive gain.

3.5. Mission Scoring

In general, standard benchmarks count successes. In the present study, FLARE measures the value of each success by defining a mission score $M(t) \in [0,1]$ that is strictly decreasing in delivery time $t \in \mathbb{N}$.

Three domain-specific instantiations specialise the FLARE framework:

(a) pharmaceutical delivery with quadratic efficacy decay,

(b) urban search-and-rescue with exponential survival decay coupled to fire proximity, and

(c) fire-perimeter surveillance with linear data-freshness decay

In domain (a), medication efficacy decays quadratically, $E(t) = \max(0, 1 - (t/T)^2)$, with the decay horizon $T = 800$ steps; the episode step budget is set to 2000 steps (effective max steps), so any successful arrival at $t \leq 800$ retains strictly positive mission value. The quadratic form reflects pharmaceutical degradation kinetics: insulin and temperature-sensitive medications retain near-full potency during an initial plateau before efficacy degrades sharply as the therapeutic window closes—a profile ill-suited to linear approximation.

In domain (b), casualty survival decays exponentially with fire-coupled amplification:

$$S_i(t) = \exp(-\lambda_{\text{eff},i} \cdot t), \lambda_{\text{eff}} = \lambda_0(1 + \kappa/\max(d_{\text{fire}}, 1)) \quad (3)$$

where λ_0 is the base decay rate—set per casualty severity to $\lambda_0 = 0.02, 0.008$, and 0.002 for severity weights 3.0, 2.0, and 1.0, respectively—and $\kappa = 5.0$ is the fire-coupling constant. Casualties near active fire expire faster, explicitly rewarding planners willing to enter dangerous zones. The fire-coupling term $\kappa/\max(d_{\text{fire}}, 1)$ captures a documented triage reality: casualties near an advancing fire front face compounding threats—heat exposure, smoke inhalation, structural instability—that accelerate physiological decline beyond the baseline trauma rate. Consequently, two casualties of identical severity at the same elapsed time may present vastly different rescue values depending solely on fire proximity at the moment of contact. The SAR mission scores and the per-family sensitivity study of Section 4.3 are all computed with this fire-coupled survival, with d_{fire} recorded at each rescue.

In domain (c), data freshness decays linearly, $F(t) = 1 - t/T$, with the same decay horizon $T = 800$ steps. A fire perimeter map becomes stale as the front advances. Linear freshness decay is standard in time-critical intelligence, surveillance, and reconnaissance (ISR) doctrine, where situational awareness degrades at a roughly constant rate as the observed scene evolves. In practice, the achievable freshness is bounded by the sensor payload and its revisit time: onboard optical and thermal imagers, their swath and data latency, and payload size, weight and power limits govern how rapidly a fire perimeter can be re-surveyed [56–59]. The surveillance score applies no explicit fire-proximity weighting, a modelling simplification we note for transparency.

All the above functions share one characteristic: they are strictly decreasing in time t , so that every additional step of detour destroys mission value.

The link between risk aversion and mission failure is now explicit. Increasing ρ produces safer paths with higher survival probability but longer travel time t . Because every mission score $M(t)$ is strictly decreasing in t , the additional travel time directly reduces mission value. This creates a fundamental trade-off:

$$\rho \uparrow (\text{more risk-averse}) \Rightarrow t \uparrow (\text{longer detour}) \Rightarrow M(t) \downarrow (\text{worse outcome})$$

Crucially, the planner receives only a target coordinate and a cost map—it has no access to POI (point of interest) locations, decay functions, or mission-level objectives. The runner orchestrates POI sequencing externally: upon task completion, it issues a new target, and upon timeout or abandonment, it redirects the planner toward the goal. The within-family ranking inversion is therefore not an artefact of mission-aware planning logic but an emergent consequence of the cost-inflation parameter ρ propagating through path geometry into arrival time and, ultimately, into mission value.

This coupling is the mechanism behind the within-family ranking inversion: within a single planner, the risk coefficient that yields the highest survival probability is penalised not by hazards but by the delivery time its detours impose—and, when inflated costs

render a task unreachable within its window, by the task value it forgoes—an inevitable consequence of any strictly time-decreasing objective.

Formally, a within-family ranking inversion is defined, for a single planner swept over its risk coefficient q , as the condition $\arg \max_q SR(q) \neq \arg \max_q M(q)$, where SR is the feasible success rate, and M is the normalised mission score. The FLARE benchmark is designed to reveal this inversion: a strictly decreasing $M(t)$ coupled to a hazard-aware cost weight $w(x)$, Equation (2), ensures that the risk coefficient maximising survival alone need not maximise mission impact.

Pareto conflict between the navigation and mission objectives. The two evaluation objectives can be written, for a planner π with risk-aversion level r (governed by q), as $J_{\text{nav}}(r) = p(r)$, the probability of reaching the goal, and $J_{\text{mis}}(r) = p(r) \cdot V(\tau(r))$, the expected mission value, where $\tau(r) = E[T \mid \text{success}]$ is the expected arrival time and V is the strictly decreasing mission-value function ($V' < 0$) of this subsection. Two structural assumptions link risk aversion to these quantities: (A1) higher risk aversion does not lower the success probability, $p'(r) \geq 0$ (wider detours avoid more hazards); and (A2) higher risk aversion does not shorten the executed path, $\tau'(r) \geq 0$ (detours cost time). Under (A1)–(A2), J_{nav} is non-decreasing in r whereas J_{mis} is the product of a non-decreasing factor and a strictly decreasing one, so the two objectives need not share a maximiser.

Differentiating, $dJ_{\text{mis}}/dr = p'(r) V(\tau(r)) + p(r) V'(\tau(r)) \tau'(r)$. Mission value therefore falls with additional risk aversion—so the more navigation-successful setting is the mission-worse one—exactly when

$$[-V'(\tau) \tau'(r)/V(\tau)] > [p'(r)/p(r)] \quad (4)$$

that is, when the relative mission value lost to delay exceeds the relative gain in success. Equation (4) defines the Pareto-conflict boundary: below it the objectives are aligned (the safest planner is also the most effective), while above it they conflict and a ranking inversion appears. Equivalently, for two risk settings $r_1 > r_2$ of a planner (so $p(r_1) > p(r_2)$) the mission ordering of the two settings inverts relative to their navigation ordering if and only if

$$\Delta p/p(r_2) < [V(\tau(r_2)) - V(\tau(r_1))]/V(\tau(r_1)) \quad (5)$$

with $\Delta p = p(r_1) - p(r_2)$: the more risk-averse setting's relative success advantage is smaller than the relative mission value it forfeits through its longer arrival time. The inversion is thus not a universal law but a regime—entered whenever the mission-value decay is steep relative to the marginal safety return—which is why, within a planner, it appears for some risk-handling paradigms and not others, and narrows to statistical noise where two settings incur near-identical delay, as presented in Section 4.3.

3.6. Reproducibility

In the present analysis, FLARE enforces 36 contracts across 15 families, verified by 197 automated tests. Key guarantees: bit-identical replay from a single RNG seed (DC-1), planner-agnostic corridor interdictions (FC-1), frozen hazard state during planning (FD-3).

Infeasible episodes are detected at reset and excluded uniformly across all planners.

3.7. Hazard Attribution via Shapley Values

Exact Shapley values [60] are a cooperative game-theory attribution measure that distributes the total success rate of a planner fairly among a set of hazard factors by averaging each factor's marginal contribution over all possible orderings of the remaining factors. For each family the Shapley value ϕ_i of each hazard coalition-member $i \in N = \{\text{fire, traffic, collapse}\}$ is computed as

$$\varphi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N|-|S|-1)!}{|N|!} \cdot [v(S \cup \{i\}) - v(S)] \quad (6)$$

where $v(S)$ is the observed success rate of that planner when only the hazards in subset S are active. A negative φ thus means the hazard degrades the planner's success rate by that amount (in success-rate units).

4. Simulation and Results

This section presents visualisations from representative simulation runs to illustrate the hazard dynamics and planner behaviour described in Section 3. A set of three cases (recent Greek wildfire events) was used to explore the effectiveness of FLARE. Hazard mechanics, perceptual differences between planners, and finally mission-value consequences are illustrated in detail. It should be stated that in our analysis, wind direction and speed are the dominant free parameters that reshape the fire front in a wildland–urban interface.

4.1. Test Cases and Scenarios

In the present study, in order to explore the effectiveness of FLARE, three selected cases were analysed that best represent different fire behaviours, impacts, and response challenges. These cases are based on Greek wildfire events with a focus on events with high severity, clear documentation, and diverse conditions [61–64], as depicted in Table 2.

- **Penteli** (Evia 2021): wind-driven fire in a wildland–urban interface; the mission is insulin delivery before efficacy degrades.
- **Piraeus** (Rhodes 2023): dense port with structural collapse; three casualties with fire-coupled survival.
- **Downtown Athens** (Evros 2023): extreme building density (0.50) with a narrow-corridor topology.

Table 2. Scenario overview. Three OSM-based scenarios inspired by recent Greek wildfires; hazard dynamics are modelled abstractions.

	Penteli	Piraeus	Downtown
Inspiring incident	Evia 2021	Rhodes 2023	Evros 2023
Mission	Insulin delivery	Search and rescue	Fire perimeter survey
Service time (hover steps at POI)	0 (fly-through)	2 steps (hover)	3 steps (hover)
Score decay $M(t)$	Quadratic	Exponential + fire	Linear
Map density (building-cell fraction)	0.18	0.29	0.50
Fire ignitions (initial count)	2	2	1
Vehicles	2 network roamers + 6 corridor patrol		
NFZ zones	0 (dynamic NFZ not activated)		
Infeasible seeds	0/30	3/30	4/30
Key challenge	Isotropic fire in low-density WUI (2 ignition points)	Collapse + segment closures	Extreme density + narrow-corridor topology

Each scenario is paired with a mission type that mirrors the dominant operational challenge of its historical precedent. Penteli maps the wildland–urban interface of the 2021 Evia megafire, where isolated settlements required time-critical medical resupply by air after road access was severed. Piraeus reproduces the dense port-adjacent topology of the

2023 Rhodes coastal fires, in which structural collapse and segment closures compounded fire hazards, stranding casualties across a fragmented urban grid. Downtown Athens reflects the 2023 Evros megafire—the largest single wildfire in EU history—where overlapping manned-aircraft corridors imposed dynamic no-fly zones on unmanned surveillance operations. These pairings ensure that the decay function, hazard dynamics, and map topology reinforce one another rather than varying independently.

4.2. Simulation Environment and Hazard Visualisation

The simulation environment is instantiated on a 500×500 grid ($N = 499$) with a cell resolution of approximately 3 m, extracted from OpenStreetMap tiles of Greek urban areas. The per-family sweep spans eight planners $\times$ three scenarios $\times$ 30 seeds $\times$ five q values (3600 episodes), of which structurally infeasible cases (7.8%) are excluded. Statistical methodology and tests are detailed in Sections 4.3 and 4.6.

Hazard-layer parameters are not free—each is tied back to a specific methodology construct. The fire radius $r_f = 30$ cells (≈ 90 m) calibrates the fire buffer defined in Section 3.3: at this radius the continuous risk layer R_f in Equation (1a) saturates for cells immediately adjacent to a burning front and decays linearly to zero at the boundary, matching the wildland–urban interface heat-exposure envelope reported in the Alexandridis cellular-automaton literature [19–21]. The traffic radius $r_t = 10$ cells (≈ 30 m) reflects the Manhattan-buffer closure radius introduced in Section 3.3 under ‘dynamic obstacles’: it is calibrated against the documented safety stand-off of low-altitude UAV corridors over moving emergency vehicles. The debris radius $r_b = 8$ cells is deliberately smaller than r_t , because collapsed-building debris is static (it cannot close an entire corridor the way a convoy can) but must still dilate sufficiently to capture rubble-field footprints. The scaling constants $c_t = 0.6$ and $c_b = 0.7$ implement the per-layer severity ordering stated in Section 3.3 (fire $>$ debris $>$ traffic at distance zero). The structural-collapse parameters $\tau_c = 80$ steps and $p_d = 0.6$ mentioned in Section 3.2 are the thresholds above which a building exposed to fire transitions to permanent debris; both were chosen so that, in a representative fire front, roughly one in two heavily exposed buildings collapses before the fire burns out, producing the monotonic degradation of the routing graph over the fire horizon. Across all three reported configurations, $\text{num_nfb_zones} = 0$; the dynamic no-fly-zone refresh module is exposed as a FLARE feature for future airspace-coordination studies and is not active in the runs reported here.

4.2.1. Penteli—Hazard Geometry and Parameter Sensitivity

The 2021 Evia fire burned 50,810 ha over eight days, driven by persistent northerly meltemi winds that elongated the active perimeter along the downwind flank. The wind-driven spread kernel [19] of Section 3.2 is available in FLARE and is illustrated on the Penteli topology in Figure 2. In more detail, this figure shows the effect of wind on fire geometry. By using isotropic fire ($\text{wind_speed} = 0$), fire expands symmetrically (see Figure 2a), whereas at wind speed 3 m/s from the west, fire elongates downwind, blocking the direct route while the upwind flank remains passable for risk-tolerant planners (see Figure 2b).

The reported Penteli benchmark uses isotropic fire ($\text{wind_speed} = 0$) so that the family comparison is not confounded by scenario-specific wind calibration. In addition, the map density is kept low (0.18) to expose an open wildland–urban interface in which the dominant obstacle is the fire front itself rather than buildings. Under this configuration, the planner-visible hazard is governed by the fire-risk radius $r_f = 30$ cells in the risk cost map of Equation (1): small changes in r_f move the effective corridor width proportionally, so Penteli is the scenario where the fire-risk radius and ignition geometry dominate the outcome. The quadratic efficacy decay $E(t)$ of the insulin mission punishes late arrivals

symmetrically around $t = T/2$, which is why static A^* is penalised even when it eventually reaches the POI—it systematically arrives on the tail of the quadratic.




Figure 2. Wind kernel illustration on the Penteli topology (illustrative seed $s = 42$, $t = 150$; the benchmark seeds are $\{0, \dots, 29\}$). (a) By using isotropic fire ($\text{wind_speed} = 0$), fire expands symmetrically. (b) Illustrative anisotropic spread with $\text{wind_speed} = 3$ m/s from the west. In both panels, light grey denotes free space, dark grey the building footprints, orange-red the burning cells and translucent grey the smoke plume; black discs with red rings are the dynamic obstacles, the blue cross (×) the UAV position at the captured step, the yellow cross the goal, and the dashed cyan line the direct start–goal corridor.

In more detail, Figure 2 contrasts the isotropic and directional regimes to document the Alexandridis-2008 [19] kernel available in FLARE for future anisotropic-spread studies. Each parameter therefore points back to a methodology section and, where possible, to an observed feature of the underlying real event—not to a tuning that was optimised for a specific planner.

Following the above, fire is transient, but its structural consequences are not. Burning buildings collapse and deposit permanent debris, monotonically eroding the available routing graph over the horizon of a single episode. In this context, Figure 3 demonstrates the cascading hazard coupling. In more detail, Figure 3a depicts that fire cell count peaks then recedes, but debris grows permanently. Four snapshots (b–e) trace the cascade from clean grid to permanent debris field, illustrating the monotonically degrading environment.

Although all planners observe the same physical hazard state, the risk coefficient q converts that shared percept into radically different effective cost surfaces; combined with each planner’s algorithm and replan trigger, these surfaces yield markedly different routes. The inflation level of each surface is governed by the risk coefficient q (Section 3.3); the underlying hazard layers that enter $R(x)$ are the fire, smoke, traffic and debris masks constructed in Section 3.3. In this context, Figure 4 holds the fire state fixed and sweeps q over $\{0, 1, 2, 5, 10\}$: at $q = 0$ the surface is uniform (risk-blind), and as q grows the fire-proximate cells inflate from passable to near-impassable while cells far from any hazard stay near unit cost. As Equation (2) is the same formula for every planner, the panels isolate the effect of q itself rather than comparing planners; each family is then swept across this q axis (Section 4.3), and how far that choice bends its executed route is shown, within one family, in Section 4.4.




Figure 3. Cascading hazard coupling on the Penteli topology (illustrative seed $s = 42$; benchmark seeds are $\{0, \dots, 29\}$). **(a)** Temporal evolution of fire and debris cell counts. **(b–e)** Four grid snapshots. In panel **(a)**, orange denotes the fire-cell count and brown the debris-cell count against the left axis, while the two faint lines give each planner's normalised distance to the goal against the right axis; the dotted vertical lines mark collapse events. In panels **(b–e)**, beige denotes free space, green the vegetation, dark grey the building footprints, orange the burning cells, black speckles the permanent debris, red-ringed discs the dynamic obstacles and the blue cross (x) the UAV position.

Same fire state ($t = 150$): the effective cost surface as the risk coefficient ρ increases




Figure 4. The effective cost surface $w(x)$, Equation (2), under a single fixed hazard state (Penteli, illustrative seed $s = 42$, $t = 150$; benchmark seeds are $\{0, \dots, 29\}$), shown as the risk coefficient ρ is swept over $\{0, 1, 2, 5, 10\}$. The surface depends on ρ alone—the same formula for every planner—so the panels isolate the effect of ρ and are not a comparison between planners. The colour bar gives the weighted traversal cost $w(x)$, from green (unit cost) through yellow to dark red (near-impassable); grey denotes the building footprints, orange squares the dynamic obstacles and the blue cross (x) the UAV position.

4.2.2. Piraeus—Coupled Hazards and Structural Collapse

The July 2023 Rhodes coastal fires triggered the largest wildfire evacuation in Greek history and caused visible structural damage in the port-adjacent zone. Our Piraeus scenario captures this signature by combining a moderate-to-high building density (0.29), the

fire–obstacle coupling with debris severity coefficient $cb = 0.7$, and the collapse cascade with threshold $\tau_c = 80$ burn-steps. The collapse cascade is enabled in all three scenario configurations, but its routing impact is most visible in Piraeus because its moderate density lets burning buildings accumulate τ_c or more burning-neighbour contacts while still leaving alternative corridors that subsequent debris can close; each collapsed building spawns permanent debris cells that inherit the fire-buffer inflation. Three casualties are placed along the A* reference corridor, and each has its own survival decay according to Equation (3).

The fire-proximity term in the denominator is what makes Piraeus the most sensitive scenario to cb and τ_c , because a single new collapse close to a casualty can double the effective decay rate observed at that victim. Incremental graph repair is well suited to these conditions, repairing the search tree as debris accumulates, whereas APF is disproportionately slowed by permanent debris cells that act as additional repulsive sources—an effect on mission timing rather than feasibility (for more details see Section 4.7).

Fire does not act in isolation: it triggers traffic closures that redirect vehicles onto alternative corridors, which themselves become new obstacles for low-altitude UAV traffic. In this context, Figure 5 illustrates the fire–obstacle coupling on the Piraeus topology: at $t = 30$ dynamic obstacles operate normally; by $t = 80$ fire triggers segment closures; and by $t = 150$ fire, closures, and diverted obstacles jointly cut the corridor, demonstrating the compound blockage mechanism.

Coupled urban dynamics: emergency vehicles, fire-induced road closures, and corridor blockage



Figure 5. Fire–obstacle coupling on the Piraeus topology (illustrative seed $s = 42$; benchmark seeds are $\{0, \dots, 29\}$). (a) $t = 30$. (b) $t = 80$. (c) $t = 150$. The basemap shows the port topology in its natural colours (blue water, grey built-up area, green vegetation); orange denotes the burning cells, dark red discs the emergency vehicles, the teal line the executed route and the blue cross (x) the UAV. The overlay in each panel reports the step index together with the fire-cell, road-closure and obstacle counts.

4.2.3. Downtown—Density, Dynamic-Obstacle Churn, and Uniform Time Pressure

The 2023 Evros megafire was the largest single wildfire ever recorded in the European Union and required sustained manned-aircraft operations that, in practice, created overlapping no-fly corridors for every other airborne platform. Downtown Athens is our hardest map in terms of static geometry (density 0.50). The dominant sensitive parameters here are the traffic-buffer radius $rt = 10$ and the debris-risk radius $rb = 8$ governing how widely dynamic obstacles and collapsed buildings inflate the risk cost map: increasing either by one cell can close the single feasible corridor for several consecutive steps, which is why Downtown records the highest infeasibility rate (4/30 seeds) of the three scenarios.

Delivery time alone determines the mission value in a time-critical operation. The three decay profiles $M(t)$ adopted for the three scenarios share this property, differing only in their analytic form (quadratic, exponential + fire coupling, linear).

Figure 6 shows the three mission-value decay profiles used by FLARE—the quadratic efficacy $E(t) = \max(0, 1 - (t/T)^2)$ for the pharma mission, the fire-coupled exponential survival $S_i(t) = \exp(-\lambda_{\text{eff},i} \cdot t)$ with $\lambda_{\text{eff}} = \lambda_0(1 + \kappa/\max(d_{\text{fire}}, 1))$, $\kappa = 5.0$ for the SAR mission, and the linear freshness $F(t) = 1 - t/T$ for surveillance. All are strictly decreasing in delivery time, so any detour destroys mission value. Markers show each planner family's median delivery step (successful episodes, pooled over the Q -sweep), coloured by risk-handling paradigm; how steeply each family's mission value falls as Q grows is quantified per family in Section 4.3.




Figure 6. The three mission-value decay profiles $M(t)$ —quadratic pharmaceutical efficacy (a), fire-coupled exponential survival (b), shown for representative fire distances d_{fire} , and linear surveillance freshness (c). In panel (b), the three line styles are the same survival law evaluated at three representative fire distances—solid for the closest casualty, dashed and dotted for progressively more distant ones—so that a shorter distance gives a steeper decay. The shaded area under the curves in panels (a) and (c) is a visual guide only and carries no additional meaning; the plotted markers give the mean delivery time of each planner family, coloured as in the legend.

4.3. Per-Family Risk-Coefficient Sweep

This section presents the paper's central result. To isolate q from planner architecture, we sweep $q \in \{0, 1, 2, 5, 10\}$ within each risk-consuming planner (three scenarios $\times$ 30 seeds), holding each planner's algorithm and replan trigger fixed. Navigation performance is the feasible success rate (SR)—the fraction of non-infeasible episodes reaching the goal within the step budget—and infeasibility is assessed at reset by a multi-depth guardrail (GC-1) that excludes episodes with no start-to-goal path at $t = 0$. A* is retained as the risk-blind static reference and excluded from the sweep: a risk-aware static A* plans off the reference corridor and thereby bypasses the corridor-aligned interdictions (FC-1), inflating navigation success without earning mission value. The eight swept planners span six risk-handling paradigm families: three soft-cost-inflation planners (Periodic, Aggressive, Incremental A*), a chance-constrained planner with a hard risk threshold [50], a CVaR/worst-case planner [51,52], a sampling-based risk-aware RRT* [53], and a risk-sensitive (entropic) planner with cost $\exp(q \cdot R)$ [54,55].

The effect of q on mission value M is strongly family-dependent, as depicted in Figures 7 and 8. It substantially erodes M in the CVaR (Spearman $\rho_s = -0.26$, $p < 0.001$; M falls from 0.85 to 0.45 as q grows from 0 to 10) and incremental soft-cost-inflation ($\rho_s = -0.25$; $0.91 \rightarrow 0.40$) families, moderately in the risk-sensitive exponential family ($0.85 \rightarrow 0.73$), and negligibly in the potential-field, hard-threshold and frequent-replan families ($|\rho_s| < 0.07$). Notably, the split runs through the soft cost-inflation family itself—the incremental planner is strongly affected while the frequently replanning Aggressive and Periodic are not—so the erosion is a property of the individual family (its cost law combined with its replan discipline) rather than of a coarse paradigm label. Navigation success, by contrast, responds mainly through a one-off gain when q becomes positive and is nearly flat thereafter (pooled $\rho_s = +0.04$ across the eight risk-parameterised families). The per-family sweep therefore separates two effects that a fixed- q cross-planner comparison conflates: the magnitude of q 's mission-value leverage, and whether the navigation- and mission-optimal q diverge within a family, as depicted in Section 4.4.

Furthermore, the effect of the risk coefficient q on each planner family (its algorithm and replan trigger held fixed) as q is swept over $\{0, 1, 2, 5, 10\}$ across the three scenarios (30 seeds per q) is presented in Table 3. Spearman(q ,SR) and Spearman(q , M) are the rank correlations of q with feasible success rate SR and with mission value M (a negative value means q erodes that metric); $q^*(SR)$ and $q^*(M)$ are the q values that maximise each; " $M: q = 0 \rightarrow 10$ " lists mission value at the lowest and highest q . When $q^*(SR) \neq q^*(M)$, a family's most navigation-successful setting is not its most mission-effective one—a within-family ranking inversion (Inv. = $\checkmark$), which occurs for Incremental A*, RRT*, Periodic and APF and is decisive for Incremental A*.




Figure 7. Per-family q -response: feasible success rate (a) and mission score M (b) versus q for the eight swept planners. Mission value collapses for the CVaR and incremental cost-inflation families, declines moderately for the risk-sensitive (exponential) family, and stays flat for the rest.




Figure 8. Spearman correlation between q and mission score M per planner family, grouped by risk-handling paradigm. q erodes M strongly for some families and negligibly for others (* $p < 0.05$).

Robustness to the mission-model parameters. To confirm the within-family inversion is not an artefact of the scoring assumptions, we re-scored all episodes over a grid of mission-model parameters—the fire-coupling constant $\kappa \in \{0, 2, 5, 10\}$, a survival-coefficient scale $\lambda \in \{0.5, 1, 2\}$, and the decay/freshness horizon $T \in \{400, 800, 1200, 2000\}$. The set of families whose navigation-optimal and mission-optimal q diverge is invariant across this entire grid—the incremental, sampling, frequent-replan, and reactive families in every case, with the incremental family's $q^*(SR) = 5$ and $q^*(M) = 0$ unchanged—so the per-family inversion is robust to the choice of scoring parameters rather than a consequence of any single calibration.

Table 3. The effect of the risk coefficient q on each planner family (its algorithm and replan trigger held fixed) as q is swept over $\{0, 1, 2, 5, 10\}$ across the three scenarios (30 seeds per q). Note: $q^*(SR)$ and $q^*(M)$ denote the value of q that maximises feasible success rate and mission value, respectively. In the “Inv.” column, $\checkmark$ marks a family whose two optima differ ($q^*(SR) \neq q^*(M)$)—a within-family ranking inversion—and — marks its absence. * $p < 0.05$.

Planner	Paradigm	Spearman (q, SR)	$q^*(SR)$	Spearman (q, M)	$q^*(M)$	M: $q = 0 \rightarrow 10$	Inv.
CVaR	Worst-case/CVaR	-0.13	1	-0.26 *	1	0.85 $\rightarrow$ 0.45	—
Incr. A*	Soft cost-inflation	+0.16	5	-0.25 *	0	0.91 $\rightarrow$ 0.40	$\checkmark$
Risk-Sensitive	Exponential/entropic	+0.03	1	-0.07	1	0.85 $\rightarrow$ 0.73	—
RRT*	Sampling	+0.08	10	-0.05	0	0.73 $\rightarrow$ 0.64	$\checkmark$
Aggressive	Soft cost-inflation	+0.08	1	-0.02	1	0.85 $\rightarrow$ 0.84	—

Chance-Constr.	Hard threshold	−0.00	1	+0.00	1	0.85 → 0.87	—
APF	Potential field	−0.01	5	+0.01	2	0.51 → 0.53	✓
Periodic	Soft cost-inflation	+0.12	2	+0.02	1	0.82 → 0.86	✓

4.4. Within-Family Ranking Inversion

Because the sweep holds each family's algorithm and replan trigger fixed, it lets us test the ranking inversion with algorithmic structure held constant: for a single family, does the q that maximises feasible success rate coincide with the one that maximises mission value? Reading $q^*(SR)$ and $q^*(M)$ from Table 3, the two optima diverge— $\arg \max_q SR(q) \neq \arg \max_q M(q)$ —for four of the eight families. The effect is decisive for the incremental cost-inflation family: its success-optimal configuration ($q = 5$, feasible SR = 64%) drives mission value down to $M = 0.41$, whereas its mission-optimal configuration ($q = 0$) attains $M = 0.91$. The sampling family (RRT*) inverts likewise, and the frequent-replan (Periodic) and reactive (APF) families invert marginally, while Aggressive, Chance-Constrained, CVaR and Risk-Sensitive place both optima at the same q . Notably, CVaR shows the strongest mission-value erosion yet no inversion, because its success- and mission-optimal q coincide at the low end—a distinction only the per-family sweep can make.

The mechanism follows from the coupling of Section 3.5. As q grows, a family inflates fire-proximate edge costs, buys a higher survival probability, and detours farther; because every mission score $M(t)$ is strictly decreasing in delivery time, the added travel time erodes mission value. For the incremental family, this trade-off is steep—each increment of q that raises feasible success also delays or forgoes task service (at $q^*(SR)$ it services the delivery POI in only 27% of feasible Penteli episodes), so its navigation-optimal configuration banks the least time-decayed value (Figure 9).




Figure 9. Within-family ranking inversion. Each family's algorithm and replan trigger are held fixed and only the risk coefficient q is swept over $\{0, 1, 2, 5, 10\}$. (a) Incremental A*—the decisive case: feasible success rate rises with q (peaking at $q^*(SR) = 5$) while mission value falls (peaking at $q^*(M) = 0$). (b) Mission value M at each family's navigation-optimal $q^*(SR)$ versus its mission-optimal $q^*(M)$.

The inversion is a regime, not a universal law (Figure 9b, Table 3). It is decisive for the incremental cost-inflation family (mission value 0.91 at $q^*(M)$ down to 0.41 at $q^*(SR)$), moderate for the sampling family (RRT*, 0.73 → 0.64), marginal for APF and Periodic, and absent for Aggressive, Chance-Constrained, CVaR and Risk-Sensitive, whose success- and mission-optimal coefficients coincide. It emerges precisely where the mission-value decay is steep relative to the marginal safety return—the Pareto boundary of Section 3.5—so it

is a property of the individual family's cost law and replan discipline, not of a paradigm label.

4.5. Qualitative Evidence

Aggregate scorecards do not by themselves reveal why a ranking inversion occurs. Overlaying the executed trajectories of a single planner across risk-coefficient settings under an identical seed exposes how the risk coefficient translates into route geometry and, subsequently, into arrival time.

Figure 10 provides spatial confirmation that, within a fixed family, the risk coefficient q governs route geometry, by holding the planner (Aggressive Replan) and the scenario fixed and varying only q over $\{0, 1, 2, 5, 10\}$ on the illustrative seed $s = 42$ (outside the benchmark range $\{0, \dots, 29\}$; the same configuration and code produce the same qualitative outcome across benchmark seeds). At $q = 0$ the planner is risk-blind and is caught by fire at $t = 266$; for every $q \geq 1$, it detours around the perimeter and reaches the goal in ≈ 522 –524 steps. The discrete flip from the direct corridor to the perimeter detour once $q \geq 1$ —with near-identical geometry thereafter—is a direct rendering of the cost-inflation law and matches this family's flat aggregate q -response (Table 3).




Figure 10. The same planner (Aggressive Replan) under an increasing risk coefficient q (Penteli, illustrative seed $s = 42$; benchmark seeds are $\{0, \dots, 29\}$). Panels, left to right: $q = 0, 1, 2, 5, 10$. At $q = 0$ the planner ignores the risk field and is caught by fire (266 steps); for every $q \geq 1$, it detours around the fire perimeter and reaches the goal (≈ 522 –524 steps). Within a fixed family, route geometry is governed by q .

4.6. Statistical Support for the Per-Family Effects

The per-family q effects are assessed by Spearman rank correlations and their significance (Table 3): the mission-value erosion is significant for the CVaR ($q_s = -0.26$, $p < 0.001$) and incremental cost-inflation ($q_s = -0.25$, $p < 0.001$) families and non-significant elsewhere, matching the family-dependent pattern of Figures 7 and 8. A Friedman test across the five q levels, blocked by family, scenario and seed ($n = 720$ blocks), confirms that q shifts mission score systematically rather than through a minority of outlier episodes ($\chi^2(4) = 229.3$, $p < 0.001$); the pooled concordance is modest (Kendall's $W \approx 0.08$) because most families are flat. Per family, the effect concentrates in the risk-sensitive ($\chi^2(4) = 91.6$), CVaR (81.9), and incremental cost-inflation (79.1) families, each $p < 0.001$. The risk-sensitive family combines a large Friedman statistic with a weak monotone correlation because its response is non-monotone—mission value peaks at $q = 1$ before declining (Figure 7b).

Interpreting the effect sizes. Cliff's δ is a non-parametric effect size in $[-1, 1]$ measuring the probability that a random episode of one configuration outscores a random episode of another; by convention $|\delta| < 0.147$ is negligible, < 0.33 small, < 0.474 medium, and ≥ 0.474 large. The Bonferroni correction multiplies each raw p -value by the number of comparisons, controlling the family-wise error rate at $\alpha = 0.05$; every effect reported as significant survives this adjustment. Applied to the two optima of each inverting family,

δ quantifies the practical size of the inversion: it is medium for the incremental cost-inflation family ($\delta = 0.44$ between M at $q^*(M)$ and M at $q^*(SR)$) and negligible for the sampling ($\delta = 0.09$), potential-field ($\delta = 0.01$) and frequent-replan ($\delta < 0.01$) families—the same ordering as the regime interpretation of Section 4.4.

4.7. Per-Family Hazard Attribution (Shapley)

Isolating dynamics layers reveals their contributions. Fire is the dominant differentiator: it reduces A^* to 0% while the adaptive families retain 39–64%. Traffic alone also eliminates A^* but has near-zero marginal impact on Periodic and Aggressive—their replanning fully mitigates obstacle blockages. Combined dynamics reduce the soft-cost-inflation family (Periodic, Aggressive, Incremental A^*) by 10–17 pp relative to fire alone, while APF is largely unaffected because its reactive field already treats traffic and debris as repulsive sources; this confirms non-additive layer interaction. Figure 11 decomposes these per-scenario contributions using exact Shapley values, confirming fire as the primary driver of feasibility loss across all three scenarios. Fire is the dominant negative contributor in every scenario; traffic is secondary and matters most for the reactive and incremental families; structural collapse contributes negligibly to success rate.




Figure 11. Scenario-separated exact Shapley attribution of the hazard layers to feasible success rate (percentage points), computed over the 2^3 hazard coalitions for each of the four classical adaptive families (Periodic, Aggressive, Incremental A^* , and APF). The collapse layer (purple in the legend) contributes close to zero percentage points in all three scenarios, so its segments are not visible at this scale; this near-zero contribution is itself a reported result (Section 5.1).

Moreover, road closures grow proportionally with fire (via the interaction engine); debris accumulates after the collapse delay ($\tau_c = 80$). Aggressive Replan’s timely detours maintain forward progress; the frequent-replan family’s fixed-interval replanning produces extended stalls whose added delay propagates into mission value through the t–M coupling.

5. Discussion

5.1. Per-Family Hazard Attribution

Hazard attribution follows the Shapley decomposition introduced in Section 3.7. The Shapley value φ_i of each hazard coalition-member $i \in N = \{\text{fire, traffic, collapse}\}$ is computed via Equation (6), where $v(S)$ is the observed success rate of the family when only the hazards in subset S are active. A negative φ thus means the hazard degrades that planner’s success rate by that amount (in success-rate units). Over the three coalitions, the Shapley decomposition confirms fire as the dominant hazard for all planners. Traffic degrades A^* ($\varphi = -0.50$) and APF ($\varphi = -0.25$), but traffic has near-zero marginal impact on

Periodic and Aggressive—their replanning absorbs obstacle blockages entirely. Collapse contributes near-zero marginal effect, subsumed by the fire that triggers it.

In fire-dominated scenarios, the within-family sweep (Table 3) shows that q is a strong mission-value lever for the CVaR and incremental cost-inflation paradigms but a negligible one for the frequent-replan and reactive families; in traffic-heavy environments, all adaptive families perform comparably.

5.2. Risk Coefficient as a Per-Family Mission-Selection Parameter

The within-family analysis establishes that no single risk setting is universally optimal, and—more sharply—that q 's leverage itself is family-dependent: the risk coefficient q serves as a meaningful, per-family mission-selection parameter.

- **Time-critical delivery (rescue, medication):** a low q , for families where the sweep shows q is an effective lever (Table 3). Minimise t to maximise $M(t)$; accept higher collision probability.
- **Collision-dominated operations (where collision yields total mission loss):** a higher q , again only where the sweep shows q moves mission value. Partial task completion retains residual value while collision yields $M = 0$. As Table 3 shows, q is a strong lever mainly for the CVaR and incremental cost-inflation families; for frequent-replan and reactive families it moves mission value only marginally, so q should be tuned where the per-family sweep shows it is effective.

This reframes planner configuration: rather than identifying a single superior algorithm at a fixed q , the appropriate risk coefficient is determined per family by the mission-specific relationship between delivery time and objective value—a relationship FLARE makes empirically quantifiable.

5.3. Limitations

Several simplifications constrain the scope of the present study.

- FLARE operates on 2D grids without altitude modelling or aerodynamic constraints.
- The planner suite comprises classical algorithms; the CA fire model provides an approximation of real wildfire dynamics.
- Three Greek urban scenarios offer focused but not exhaustive geographic coverage.
- The risk coefficient is swept over a discrete grid $q \in \{0, 1, 2, 5, 10\}$ and held constant within each episode; online or state-dependent adaptation of q is not studied.
- Because conventional, navigation-centric benchmarks do not score the time-dependent value of mission completion, the within-family ranking inversion is not visible under standard evaluation protocols; FLARE surfaces it by sweeping q within each family against a time-decaying mission value. The inversion is a per-family regime rather than a universal law: it appears in four of the eight families studied and is absent in the rest.
- The surveillance freshness model abstracts sensor revisit time, data latency, and payload size, weight and power constraints [56–59], and applies no fire-proximity weighting.

The above simplifications are deliberate: the contribution of the present study is the evaluation paradigm and the ranking-inversion finding, not the planners themselves.

Finally, the benchmark operates on a 2-D grid; a third (altitude) dimension would make the risk field $R(x, y, z)$. Vertical separation attenuates horizontal ground-hazard risk (fire plumes, debris, moving vehicles) but introduces altitude-dependent terms—smoke-column height, manned-aircraft corridors and dynamic no-fly volumes, and wind shear—so extending FLARE to 3-D airspace, where altitude reshapes the risk-cost map, is a priority direction.

6. Conclusions and Future Steps

This paper presented FLARE, a deterministic benchmark for evaluating UAV planner performance under time-coupled mission objectives. Its central instrument is a within-family sweep of the risk coefficient q , with each planner's algorithm held fixed, across 450 base episodes on three Greek wildfire scenarios. The sweep shows that q 's leverage on mission value is family-dependent—decisive for the incremental cost-inflation family (mission value 0.41 at its success-optimal q versus 0.91 at its mission-optimal q), strong for the worst-case (CVaR) family, and negligible for the reactive, hard-threshold and frequent-replan families—and that in four of the eight families the q maximising navigation success is not the q maximising mission value (a within-family ranking inversion), while in the rest the two optima coincide.

Because the algorithm is fixed across each sweep, this divergence is attributable to q alone rather than to differences between planner architectures; it is a regime, not a universal law. The principal implication is methodological: planner evaluation should incorporate time-dependent mission-value functions rather than binary completion metrics, and the risk coefficient q provides a meaningful, per-family parameter for mission-specific configuration.

Future work will extend FLARE in three directions. First, the discrete 2D grid will be generalised to a 3D airspace with altitude-coupled hazard perception—and ultimately to continuous-motion kinematics—so that smoke column height and no-fly zones of finite vertical extent can be represented explicitly. Second, learning-based planners (reinforcement-learning policies and graph neural-network cost regressors) will be added to the planner set so that the risk-coefficient calibration protocol can be applied uniformly to data-driven agents. Third, the fire cellular automaton will be upgraded from a Moore-neighbourhood Alexandridis model to a fuel-coupled Rothmel variant, enabling explicit calibration against historical EFFIS/Copernicus fire perimeters.

Author Contributions: Conceptualization, K.Z. and I.P.; methodology, K.Z.; software, K.Z.; validation, K.Z. and I.P.; formal analysis, K.Z.; investigation, K.Z.; resources, K.Z.; data curation, K.Z.; writing—original draft preparation, K.Z.; writing—review and editing, K.Z. and I.P.; visualisation, K.Z.; supervision, I.P.; project administration, K.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The benchmark code, scenario configurations, and full experimental results are publicly available at <https://github.com/zervakisai/uavplan> (accessed on 15 June 2026). Map data are derived from OpenStreetMap under the Open Database License (ODbL).

Acknowledgments: Map data © OpenStreetMap contributors (ODbL). Scenarios inspired by the 2021 Evia wildfire [61,62], 2023 Rhodes fires [63], and 2023 Evros megafire [3,64]. Data from EFFIS [61,64] and Copernicus EMS [3,62,63]. During the preparation of this manuscript, the authors used Perplexity AI (web application; <https://www.perplexity.ai/>—accessed December 2025–January 2026) for the purpose of literature research and source discovery. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

UAV	Unmanned Aerial Vehicle
FLARE	Fire-Landscape Adaptive Risk Evaluation
SR	Success Rate
MS	Normalised Mission Score
IQM	Interquartile Mean
CI	Confidence Interval
CA	Cellular Automaton
APF	Artificial Potential Field
OSM	OpenStreetMap
NFZ	No-Fly Zone
RNG	Random Number Generator
WUI	Wildland–Urban Interface
SAR	Search and Rescue
CVaR	Conditional Value-at-Risk
RRT*	Rapidly exploring Random Tree (optimal variant)

References

1. Erdelj, M.; Natalizio, E.; Chowdhury, K.R.; Akyildiz, I.F. Help from the Sky: Leveraging UAVs for Disaster Management. *IEEE Pervasive Comput.* **2017**, *16*, 24–32. <https://doi.org/10.1109/MPRV.2017.11>.
2. Yücesoy, G.; Balcik, B.; Coban, E. The Role of drones in disaster response: A literature review of operations research applications. *Int. Trans. Oper. Res.* **2025**, *32*, 545–589. <https://doi.org/10.1111/itor.13484>.
3. Copernicus Emergency Management Service. EMSN166: Wildfire in Evros, Greece—Risk and Recovery Mapping. Available online: <https://mapping.emergency.copernicus.eu/activations/EMSN166/> (accessed on 30 July 2026).
4. Li, J.; Pei, Y.; Zhao, S.; Xiao, R.; Sang, X.; Zhang, C. A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management. *Inf. Fusion* **2024**, *108*, 102369. <https://doi.org/10.1016/j.inffus.2024.102369>.
5. Meng, W.; Zhang, X.; Zhou, L.; Guo, H.; Hu, X. Advances in UAV Path Planning: A Comprehensive Review of Methods, Challenges, and Future Directions. *Drones* **2025**, *9*, 376. <https://doi.org/10.3390/drones9050376>.
6. Hart, P.E.; Nilsson, N.J.; Raphael, B. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Trans. Syst. Sci. Cybern.* **1968**, *4*, 100–107. <https://doi.org/10.1109/TSSC.1968.300136>.
7. Koenig, S.; Likhachev, M. D* Lite. In Proceedings of the 18th AAAI Conference on Artificial Intelligence, Edmonton, AB, Canada, 28 July–1 August 2002; pp. 476–483.
8. Khatib, O. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *Int. J. Robot. Res.* **1986**, *5*, 90–98. <https://doi.org/10.1177/027836498600500106>.
9. Aggarwal, S.; Kumar, N. Path Planning Techniques for Unmanned Aerial Vehicles: A Review, Solutions, and Challenges. *Comput. Commun.* **2020**, *149*, 270–299. <https://doi.org/10.1016/j.comcom.2019.10.014>.
10. Chen, H.; Lin, Y.; Fu, M.; Yao, L.; Sheng, M. A Survey on Reinforcement Learning Methods for UAV Systems. *ACM Comput. Surv.* **2025**, *58*, 1–37. <https://doi.org/10.1145/3769426>.
11. Bouzid, Y.; Bestaoui Sebbane, Y.; Mohammedi, A. UAV Path Planning Trends from 2000 to 2024: A Bibliometric Analysis and Visualization. *Drones* **2025**, *9*, 128. <https://doi.org/10.3390/drones9020128>.
12. Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; Meger, D. Deep Reinforcement Learning That Matters. In Proceedings of the 32nd AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
13. Zhou, W.; Wang, K.; Wang, Y.; Qu, X. A risk-based unmanned aerial vehicle path planning scheme for complex air-ground environments. *Risk Anal.* **2025**, *45*, 321–342. <https://doi.org/10.1111/risa.17685>.
14. Melissaris, C.; Zeh, T.; Lindner, M.; Fricke, H. Risk-aware UAV trajectory optimization using open urban GIS data and target level of safety constraints. *Drones* **2025**, *9*, 666. <https://doi.org/10.3390/drones9100666>.
15. Karaağaç, C.; Yilmaz, N.; Hançerlioğulları, A. Risk-Aware Enabled Path Planning for Drones Flight in Unknown Environment. *J. Intell. Robot. Syst.* **2025**, *111*, 56. <https://doi.org/10.1007/s10846-025-02243-2>.
16. Chen, X.; Zhang, F.; Liu, Y.; Zhang, L. Probabilistic Path Planning for UAVs in Forest Fire Monitoring: Enhancing Patrol Efficiency through Risk Assessment. *Fire* **2024**, *7*, 254. <https://doi.org/10.3390/fire7070254>.

17. Roberts, N.B.; Ager, E.; Leith, T.; Lott, I.; Mason-Maready, M.; Nix, T.; Gottula, A.; Hunt, N.; Brent, C. Current Summary of the Evidence in Drone-Based Emergency Medical Services Care. *Resusc. Plus* **2023**, *13*, 100347. <https://doi.org/10.1016/j.resplu.2022.100347>.
18. Thompson, M.P.; Calkin, D.E.; Herynk, J.; McHugh, C.W.; Short, K.C. Risk Management and Analytics in Wildfire Response. *Curr. For. Rep.* **2019**, *5*, 226–239. <https://doi.org/10.1007/s40725-019-00101-7>.
19. Alexandridis, A.; Vakalis, D.; Siettos, C.I.; Bafas, G.V. A Cellular Automata Model for Forest Fire Spread Prediction: The Case of the Wildfire that Swept Through Spetses Island in 1990. *Appl. Math. Comput.* **2008**, *204*, 191–201. <https://doi.org/10.1016/j.amc.2008.06.046>.
20. Sahila, A.; Canfora, B.; Canzaniello, M.; Perello, N.; Trucchia, A.; Pagnini, G. Statistical dynamics of wildfire burned area from cellular-automata simulators. *Sci. Rep.* **2026**, *16*, 998. <https://doi.org/10.1038/s41598-025-30851-3>.
21. Sahila, A.; Fiorese, A.; Perello, N.; Trucchia, A.; Pagnini, G. Patterns of burned area by a cellular-automata fire simulator: The role of microscale wind field. *Commun. Nonlinear Sci. Numer. Simul.* **2025**, *151*, 109026. <https://doi.org/10.1016/j.cnsns.2025.109026>.
22. Wang, Z.; Zhao, M.; Chen, G.; Li, J. Multi-Target Firefighting Task Planning Strategy for Multiple UAVs Under Dynamic Forest Fire Environment. *Fire* **2025**, *8*, 61. <https://doi.org/10.3390/fire8020061>.
23. Liu, Y.; Wang, Z.; Chen, X.; Zhang, W. Multiple UAV Swarms Collaborative Firefighting Strategy Considering Forest Fire Spread and Resource Constraints. *Drones* **2025**, *9*, 17. <https://doi.org/10.3390/drones9010017>.
24. Quero, C.O.; Martinez-Carranza, J. Unmanned aerial systems in search and rescue: A global perspective on current challenges and future applications. *Int. J. Disaster Risk Reduct.* **2025**, *118*, 105199. <https://doi.org/10.1016/j.ijdrr.2025.105199>.
25. Habibi, S.; Ivaki, N.; Barata, J. A systematic literature review of unmanned aerial vehicles for healthcare and emergency services. *arXiv* **2025**, arXiv:2504.08834.
26. Al-Atroshi, C.; Sarkar, N.I.; Lutui, R. A survey on multi-UAV path planning: Classification, algorithms, open research problems, and future directions. *Drones* **2025**, *9*, 263. <https://doi.org/10.3390/drones9040263>.
27. Dradoun, A.; Khelassi, A.; Lachekhab, F. Intelligent path planning algorithms for UAVs: Classification, complexity analysis, hybrid ablation insights, and future directions. *Adv. Mech. Eng.* **2025**, *17*, 1–26. <https://doi.org/10.1177/16878132251355020>.
28. Jiang, Y.; Xu, X.-X.; Zheng, M.-Y. Evolutionary computation for unmanned aerial vehicle path planning: A survey. *Artif. Intell. Rev.* **2024**, *57*, 267. <https://doi.org/10.1007/s10462-024-10913-0>.
29. Primatesta, S.; Rizzo, A.; la Cour-Harbo, A. A risk-aware path planning strategy for UAVs in urban environments. *J. Intell. Robot. Syst.* **2019**, *95*, 629–643. <https://doi.org/10.1007/s10846-018-0924-3>.
30. Tang, H.; Zhu, Q.; Qin, B.; Song, R. UAV path planning based on third-party risk modeling. *Sci. Rep.* **2023**, *13*, 22259. <https://doi.org/10.1038/s41598-023-49396-4>.
31. Agarwal, R.; Schwarzer, M.; Castro, P.S.; Courville, A.; Bellemare, M.G. Deep Reinforcement Learning at the Edge of the Statistical Precipice. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021), Online, 6–14 December 2021*; Curran Associates Inc.: Red Hook, NY, USA, 2021; Volume 34.
32. Demšar, J. Statistical Comparisons of Classifiers over Multiple Data Sets. *J. Mach. Learn. Res.* **2006**, *7*, 1–30.
33. Dolan, E.D.; Moré, J.J. Benchmarking Optimization Software with Performance Profiles. *Math. Program.* **2002**, *91*, 201–213. <https://doi.org/10.1007/s101070100263>.
34. Bonsignorio, F.; del Pobil, A.P.; Zereik, E. Towards reproducible robotics research. *Nat. Mach. Intell.* **2025**, *7*, 15–22. <https://doi.org/10.1038/s42256-025-01114-7>.
35. Zhao, Y.; Zheng, Z.; Liu, Y. Survey on Computational-Intelligence-Based UAV Path Planning. *Knowl.-Based Syst.* **2018**, *158*, 54–64. <https://doi.org/10.1016/j.knsys.2018.05.033>.
36. Boutilier, J.J.; Brooks, S.C.; Janmohamed, A.; Byers, A.; Buick, J.E.; Zhan, C.; Schoellig, A.P.; Cheskes, S.; Morrison, L.J.; Chan, T.C.Y. Optimizing a Drone Network to Deliver Automated External Defibrillators. *Circulation* **2017**, *135*, 2454–2465. <https://doi.org/10.1161/CIRCULATIONAHA.116.026318>.
37. Claesson, A.; Bäckman, A.; Ringh, M.; Svensson, L.; Nordberg, P.; Djärv, T.; Hollenberg, J. Time to Delivery of an Automated External Defibrillator Using a Drone for Simulated Out-of-Hospital Cardiac Arrests vs Emergency Medical Services. *JAMA* **2017**, *317*, 2332–2334. <https://doi.org/10.1001/jama.2017.3957>.
38. Nisingizwe, M.P.; Ndishimye, P.; Swaibu, K.; Nshimiyimana, L.; Karame, P.; Dushimiyimana, V.; Musabyimana, J.P.; Musanabaganwa, C.; Nsanzimana, S.; Law, M.R. Effect of unmanned aerial vehicle (drone) delivery on blood product delivery time and wastage in Rwanda: A retrospective, cross-sectional study and time series analysis. *Lancet Glob. Health* **2022**, *10*, e564–e569. [https://doi.org/10.1016/S2214-109X\(22\)00048-1](https://doi.org/10.1016/S2214-109X(22)00048-1).

39. Vansteenwegen, P.; Souffriau, W.; Van Oudheusden, D. The orienteering problem: A survey. *Eur. J. Oper. Res.* **2011**, *209*, 1–10. <https://doi.org/10.1016/j.ejor.2010.03.045>.
40. Gunawan, A.; Lau, H.C.; Vansteenwegen, P. Orienteering Problem: A survey of recent variants, solution approaches and applications. *Eur. J. Oper. Res.* **2016**, *255*, 315–332. <https://doi.org/10.1016/j.ejor.2016.04.059>.
41. Shah, S.; Dey, D.; Lovett, C.; Kapoor, A. AirSim: High-Fidelity Visual and Physical Simulation for Autonomous Vehicles. In *Field and Service Robotics*; Springer: Cham, Switzerland, 2018; pp. 621–635. https://doi.org/10.1007/978-3-319-67361-5_40.
42. Song, Y.; Naji, S.; Kaufmann, E.; Loquercio, A.; Scaramuzza, D. Flightmare: A Flexible Quadrotor Simulator. In *Proceedings of the Conference on Robot Learning (CoRL), Virtual, 16–18 November 2020*; PMLR: New York, NY, USA, 2020; Volume 155.
43. Panerati, J.; Zheng, H.; Zhou, S.; Xu, J.; Prorok, A.; Schoellig, A.P. Learning to Fly—A Gym Environment with PyBullet Physics for Reinforcement Learning of Multi-Agent Quadcopter Control. In *Proceedings of the IEEE/RSJ IROS, Prague, Czech Republic, 27 September–1 October 2021*; pp. 7512–7519. <https://doi.org/10.1109/IROS51168.2021.9635857>.
44. Seraj, E.; Wu, X.; Gombolay, M. FireCommander: An Interactive, Probabilistic Multi-Agent Environment for Heterogeneous Robot Teams. *arXiv* **2020**, arXiv:2011.00165.
45. Hyun, J.; Waytowich, N.R.; Chen, B. CREW-Wildfire: Benchmarking Agentic Multi-Agent Collaborations at Scale. *arXiv* **2025**, arXiv:2507.05178.
46. Finney, M.A. *FARSITE: Fire Area Simulator—Model Development and Evaluation*; Research Paper RMRS-RP-4; USDA Forest Service, Rocky Mountain Research Station: Ogden, UT, USA, 1998.
47. Kelly, J.R.; Chaudhry, U.B. Mission-Driven UAV Path Selection: Post Hoc Cost Evaluation of Deterministic and Sampling Approaches. *Drones* **2026**, *10*, 152. <https://doi.org/10.3390/drones10020152>.
48. Towers, M.; Kwiatkowski, A.; Terry, J.K.; Balis, J.U.; De Cola, G.; Deleu, T.; Goulão, M.; Kallinteris, A.; Krimmel, M.; Arjun, K.G.; et al. Gymnasium: A Standard Interface for Reinforcement Learning Environments. In *Proceedings of the Advances in Neural Information Processing Systems, San Diego, CA, USA, 2–7 December 2025*; Curran Associates, Inc.: Red Hook, NY, USA, 2025; Volume 38.
49. Boeing, G. OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks. *Comput. Environ. Urban Syst.* **2017**, *65*, 126–139. <https://doi.org/10.1016/j.compenvurbsys.2017.05.004>.
50. Blackmore, L.; Ono, M.; Williams, B.C. Chance-Constrained Optimal Path Planning with Obstacles. *IEEE Trans. Robot.* **2011**, *27*, 1080–1094. <https://doi.org/10.1109/TRO.2011.2161160>.
51. Rockafellar, R.T.; Uryasev, S. Optimization of Conditional Value-at-Risk. *J. Risk* **2000**, *2*, 21–41. <https://doi.org/10.21314/JOR.2000.038>.
52. Majumdar, A.; Pavone, M. How Should a Robot Assess Risk? Towards an Axiomatic Theory of Risk in Robotics. In *Robotics Research (ISRR 2017)*; Springer: Cham, Switzerland, 2020; pp. 75–84. https://doi.org/10.1007/978-3-030-28619-4_10.
53. Karaman, S.; Frazzoli, E. Sampling-based algorithms for optimal motion planning. *Int. J. Robot. Res.* **2011**, *30*, 846–894. <https://doi.org/10.1177/0278364911406761>.
54. Jacobson, D.H. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Autom. Control* **1973**, *18*, 124–131. <https://doi.org/10.1109/TAC.1973.1100265>.
55. Whittle, P. Risk-sensitive linear/quadratic/Gaussian control. *Adv. Appl. Probab.* **1981**, *13*, 764–777. <https://doi.org/10.2307/1426972>.
56. Allison, R.S.; Johnston, J.M.; Craig, G.; Jennings, S. Airborne Optical and Thermal Remote Sensing for Wildfire Detection and Monitoring. *Sensors* **2016**, *16*, 1310. <https://doi.org/10.3390/s16081310>.
57. Thornberry, T.D.; Gao, R.-S.; Ciciora, S.J.; Watts, L.A.; McLaughlin, R.J.; Leonardi, A.; Rosenlof, K.H.; Argrow, B.M.; Elston, J.S.; Stachura, M.; et al. A Lightweight Remote Sensing Payload for Wildfire Detection and Fire Radiative Power Measurements. *Sensors* **2023**, *23*, 3514. <https://doi.org/10.3390/s23073514>.
58. Schuyler, T.J.; Guzman, M.I. Unmanned Aerial Systems for Monitoring Trace Tropospheric Gases. *Atmosphere* **2017**, *8*, 206. <https://doi.org/10.3390/atmos8100206>.
59. Guzman, M.I. Atmospheric Measurements with Unmanned Aerial Systems (UAS). *Atmosphere* **2020**, *11*, 1208. <https://doi.org/10.3390/atmos11111208>.
60. Shapley, L.S. A Value for n-Person Games. In *Contributions to the Theory of Games II*; Kuhn, H.; Tucker, A., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; pp. 307–317.
61. San-Miguel-Ayanz, J.; Durrant, T.; Boca, R.; Maiani, P.; Libertà, G.; Artés-Vivancos, T.; Jacome Felix Oom, D.; Branco, A.; De Rigo, D.; Ferrari, D.; et al. *Forest Fires in Europe, Middle East and North Africa 2021*; Technical report; EUR 31269 EN; JRC130846; European Commission, Joint Research Centre: Luxembourg, 2022; ISBN 978-92-76-58585-5. <https://doi.org/10.2760/34094>.

62. Copernicus Emergency Management Service. EMSR527: Wildfires in Evia and Attica, Greece, August 2021—Rapid Mapping. Available online: <https://mapping.emergency.copernicus.eu/activations/EMSR527/> (accessed on 30 July 2026).
63. Copernicus Emergency Management Service. EMSR675: Wildfire on Rhodes island, Greece, 2023—Rapid Mapping. Available online: <https://mapping.emergency.copernicus.eu/activations/EMSR675/> (accessed on 30 July 2026).
64. San-Miguel-Ayanz, J.; Durrant, T.; Boca, R.; Maianti, P.; Libertà, G.; Jacome Felix Oom, D.; Branco, A.; De Rigo, D.; Suarez-Moreno, M.; Ferrari, D.; et al. *Forest Fires in Europe, Middle East and North Africa 2023*; Technical report; JRC139704; European Commission, Joint Research Centre: Luxembourg, 2024. <https://doi.org/10.2760/8027062>.

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