AI Engineering
https://www.sciltp.com/journals/aieng



Article

Real-Time Fire Escape Path Guidance System through YOLO11 and Dynamic A* Algorithm

Jun Ang [1], Lei Guo [2], Xiangyu Liu [1], Tianyu Shi [1,*], Lincong Zhang [1] and Qiang Fu [3]

[1] School of Information Science and Engineering, Shenyang Ligong University, Shenyang 110182, China
[2] School of Computer Science and Engineering, Northeastern University, Shenyang 110169, China
[3] Nanjing Automation Research Institute (NARI) Technology Company Limited, Nanjing 210023, China
* Correspondence: ty.shi@sylu.edu.cn

How To Cite: Ang, J.; Guo, L.; Liu, X.; et al. Real-Time Fire Escape Path Guidance System through YOLO11 and Dynamic A* Algorithm. *AI Engineering* **2026**, *2*(2), 22. https://doi.org/10.53941/aieng.2026.100022

Received: 17 June 2026
Revised: 23 October 2026
Accepted: 26 October 2026
Published: 11 November 2026

Abstract: We present a closed-loop fire escape guidance system that integrates YOLO11 flame detection, map-aligned hazard mapping, and event-triggered A* replanning. Corridor-camera detections are converted into grid cells, expanded by a 3 × 3 safety buffer, and updated by a time-dependent four-neighbor fire-spread model. A* recomputes routes to the nearest stairwell when a new flame appears, fire spreads, a restored cell reopens, or the start position changes. The system displays the live hazard map and route and transmits alarms through the ServerChan Turbo API to WeChat (version 3.9.10.19, Tencent, Shenzhen, China). Each alarm contains the detection time, start position, path length, and flame count. On the fixed validation split, the deployment YOLO11m checkpoint achieves 77.34% mAP@0.5; on a duplicate-aware group-disjoint test set, the final model achieves 72.36%. A single A* call averages 0.407 ± 0.011 ms on an Intel Core i5-14600KF CPU (Intel Corporation, Santa Clara, CA, USA). In a controlled 60 s scenario with 3 s updates, the system finds a fire-free, buffer-clear route at 30.0% of replanning events, and the common-timeline evacuation run reaches the stair region in 15.83 s. Grid maps from 21 × 21 to 81 × 81 are processed successfully.

Keywords: flame detection; YOLO11; A* algorithm; path planning

1. Introduction

Building fires threaten urban safety, while fixed evacuation signs cannot adapt to fire location or spread. Vision-based fire management links sensing, localization, decision support and response [1], and feature-fusion methods improve small-target representation [2]. Event-driven control provides an efficient mechanism for organizing repeated state updates [3]. Recent integrated work has coupled semi-occluded flame detection with evacuation-route generation [4]. A* provides minimum-cost search on a weighted graph [5]. This study integrates these elements into a closed-loop workflow in which YOLO11 detections update a grid hazard map and trigger route recomputation.

This work provides an engineering integration and system-level evaluation. First, YOLO11 [6] detection boxes are mapped to grid hazard cells and expanded by a 3 × 3 safety buffer. Second, time-dependent fire growth updates the same grid; new-flame detections, fire spread, restoration, and start-position changes trigger A* replanning and refresh the guidance display and alarms. Third, validation combines a duplicate-aware group-disjoint detector test, planner benchmarks, shared-timeline ablation, and a common-timeline evacuation run. Recent building-fire studies provide complementary unsafe-area prediction [7] and tenability-aware routing [8]. A* remains an established planner [5]; “dynamic” denotes event-triggered re-invocation.

Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
Publisher’s Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

2. Overall System Design

System Architecture

The system comprises camera acquisition, YOLO11 detection, grid update, A* planning, visualization, and WeChat alarm modules. Confirmed flame boxes are mapped to grid cells and expanded by a 3 × 3 safety buffer. A time-dependent four-neighbor model updates fire zones every 3 s, and A* replans to the nearest stairwell. Matplotlib (version 3.11.1) displays the map, fire zones, route, and start position; mouse clicks update that position, and alarms report the detection time, start position, path length, and flame count. New-flame detections, fire spread, restoration, and start-position changes update the planning state.

3. Flame Detection Module

3.1. Dataset and Preprocessing

The curated dataset contains 5985 images: 3311 originals and 2674 augmented copies. Exact and near-duplicate images (dHash ≤ 4) were grouped before splitting, producing 2180 groups. Of the 5985 files, 800 augmented copies belonging to validation or test groups were withheld from the evaluation lists so that validation and test contain originals only. The final training, validation, and independent test lists therefore contain 4191/497/497 images (5185 in total) and 1526/327/327 groups, with zero exact or dHash ≤ 4 pairs shared with training. All images use YOLO-format manual annotations. Online augmentation includes horizontal flipping, saturation and brightness adjustment, and Mosaic.

3.2. Model Training

YOLO11m was trained for 20 epochs. Precision, recall, mAP@0.5, and mAP@0.5:0.95 increased through epoch 15 and stabilized through epoch 20 (Figure 1). The fixed validation split was used for model development, and Section 5.5 reports the group-disjoint independent test.

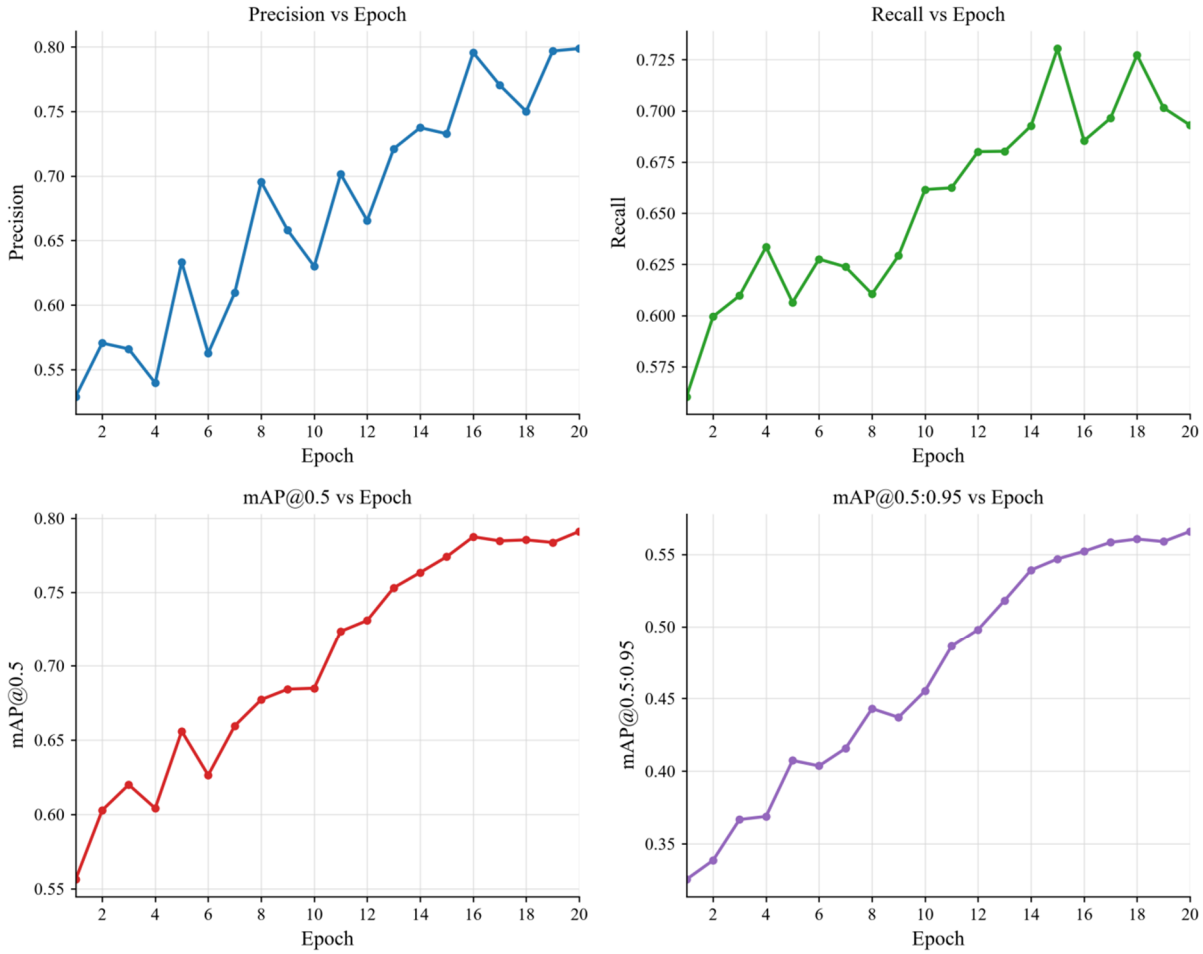


Figure 1. YOLO11m training epochs.

3.3. Detector Comparison and Model Selection

Detection performance is reported with precision, recall, mAP@0.5, mAP@0.5:0.95, parameter count, model size, and single-image latency. Table 1 reports the two YOLO11 deployment checkpoints on the fixed validation split. Table 2 reports a separate controlled m-level comparison of YOLO11m, YOLOv8m and YOLOv10m under a common training and validation protocol. The two tables serve distinct evaluation purposes and are not pooled into a single model ranking.

Table 1. YOLO11m and YOLO11n deployment-checkpoint results on the fixed validation split.

	Training Time (s)	Precision (%)	Recall (%)	mAP@0.5 (%)	mAP@0.5:0.95 (%)
YOLO11m	1531.68	73.39	71.77	77.34	55.43
YOLO11n	446.493	72.47	67.60	75.93	52.68

Table 2. Controlled m-level detector comparison for architecture selection under a common training and validation protocol.

Model	Precision (%)	Recall (%)	mAP@0.5 (%)	mAP@0.5:0.95 (%)	Parameters (M)/Size (MB)	Median Latency (ms)
YOLOv8m	76.22	69.22	77.87	55.35	25.84/49.59	15.40
YOLOv10m	75.50	70.21	78.00	56.27	15.31/31.89	6.86
YOLO11m	79.75	69.00	79.12	56.49	20.03/38.60	8.01

YOLO11m improves mAP@0.5 by 1.41 percentage points and mAP@0.5:0.95 by 2.75 percentage points over YOLO11n. Its training time is 1531.68 s versus 446.493 s for YOLO11n, so YOLO11m is used for accuracy-oriented deployment and YOLO11n for resource-constrained deployment. These are deployment-checkpoint validation results; the controlled architecture comparison is reported separately in Table 2, and the group-disjoint independent estimate is reported in Section 5.5.

YOLOv8m, YOLOv10m and YOLO11m were trained for 20 epochs on the same dataset and fixed split with 512 × 512 input, AdamW, an initial learning rate of 0.001, batch size 8 and a fixed seed. Table 2 reports the resulting accuracy and computational characteristics. Tables 1 and 2 correspond to separate training runs and evaluation scopes, so their percentage values are not interchangeable.

Latency is the median single-image prediction time over 100 measured calls after 20 warm-up calls on an NVIDIA RTX 4060 (NVIDIA Corporation, Santa Clara, CA, USA). Each call uses a 512 × 512 input and confidence 0.25; preprocessing, inference, and post-processing are included, and disk I/O is excluded.

YOLO11m achieved the highest precision (79.75%), mAP@0.5 (79.12%), and mAP@0.5:0.95 (56.49%) in the controlled m-level comparison. YOLOv10m achieved 70.21% recall, 15.31 M parameters, 31.89 MB model size, and 6.86 ms latency. YOLO11m uses 20.03 M parameters, 38.60 MB storage, and 8.01 ms latency, providing the accuracy-oriented choice; YOLOv10m provides a lower-cost alternative. The 79.12% value is used only for detector-architecture selection and is not an independent-test estimate; deployment-checkpoint validation and group-disjoint independent results are reported in Table 1 and Section 5.5, respectively.

3.4. Flame Position Conversion in the Map-Aligned Virtual-Camera Frame

Let (c_x, c_y) be a confirmed detection center and $W \times H$ the frame size. The normalized image coordinates are

$$u = \frac{c_x}{W}, v = \frac{c_y}{H} \tag{1}$$

For a fixed camera, the planar homography H maps the normalized image point to normalized map coordinates:

$$[X'\ Y'\ W']^T = H[u\ v\ 1]^T, X = \frac{X'}{W'}, Y = \frac{Y'}{W'} \tag{2}$$

N_c and N_r are the numbers of grid columns and rows. The map coordinates are converted to a 1-based cell by

$$x = \mathrm{clamp}(\mathrm{floor}(X\ N_c) + 1, 1, N_c), y = \mathrm{clamp}(\mathrm{floor}(Y\ N_r) + 1, 1, N_r) \tag{3}$$

With $H = I$ for the map-aligned camera, Equations (1)–(3) provide the common quantization path used by the calibrated-camera adapter. Coordinate checks covered all 1681 cell centers of the 41 × 41 map, four image corners, the image center and boundary probes; every point returned the intended in-range cell. The homography adapter was checked with eight held-out control points in a controlled projective test, with a maximum reprojection error below 2.2×10^{-13} map-coordinate units.

Each confirmed detection marks its corresponding cell and the full 3 × 3 neighborhood as a fire zone. The 30-frame confirmation rule requires sustained detections before an alarm, and the counter resets when a frame has

no flame target. At 30 fps, the rule provides a 1.0 s observation window and filters brief visual disturbances. Fire-position conversion, map update, and WeChat alarm workflows start after confirmation.

The confirmation rule was evaluated with YOLO11m on a 302-frame, 30-fps flame recording at 512 × 512 input, confidence 0.6, and IoU 0.7. The first uninterrupted above-threshold segment began at frame 155 (5.17 s) and contained 31 frames. The 20-frame and 30-frame rules were confirmed at 5.83 s and 6.17 s, corresponding to delays of 0.67 s and 1.00 s; the 40-frame rule requires a 1.33 s observation window and was not satisfied within this segment. The 30-frame rule therefore provides a 1.0 s evidence window and confirms the tested flame event. All planner times in Sections 5.1–5.4 measure a single A* call.

4. Dynamic Path Planning Module

4.1. Environmental Modeling and Target Setting

The environment is represented by a 41 × 41 grid (Figure 2), where 0, 1, 2, and 3 denote corridors, walls, stairwells, and fire zones, respectively. In Figure 2, white, black, and red cells represent corridors, walls, and fire zones, respectively. All stairwell cells are valid targets, and fire zones toggle between active and inactive states to support efficient obstacle checks and route updates.

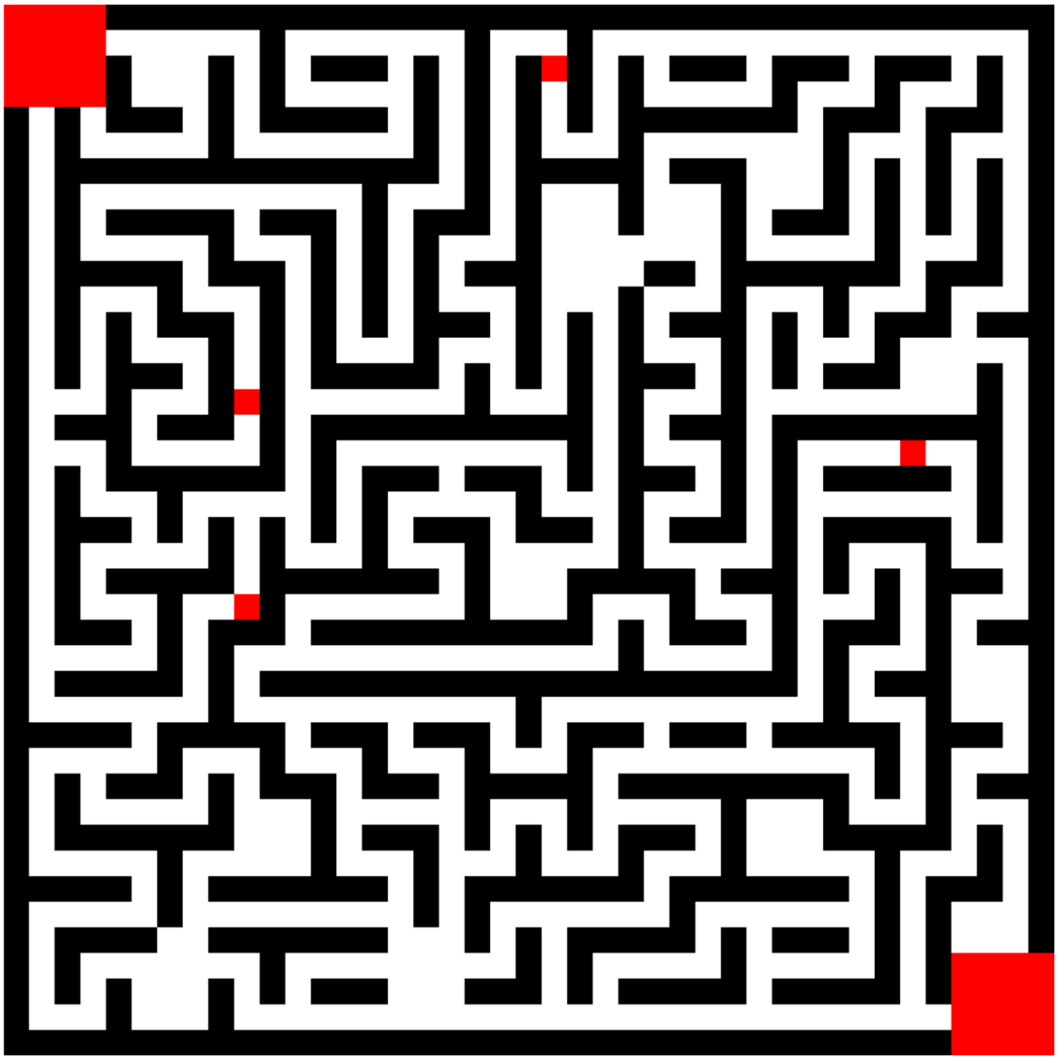


Figure 2. Simulated map.

4.2. A Algorithm Implementation*

A* is an established graph-search planner. It minimizes $f(n) = g(n) + h(n)$, where $g(n)$ is the accumulated path cost from the start and $h(n)$ is the Manhattan distance to the target. Four-neighbor movement makes the Manhattan heuristic consistent and returns a minimum-cost path on the tested finite grid whenever one exists.

$$f(n) = g(n) + h(n), h(n) = |x - x_\text{goal}| + |y - y_\text{goal}| \quad (4)$$

The open set uses a binary min-heap, and the search expands four neighbors with O(VlogV) time and O(V) memory for V traversable cells. Confirmed flames and their 3 × 3 buffer are non-walkable, while wall and fire cells are excluded from expansion. New-flame, spread, restoration, and start-position changes update the grid and invoke A* replanning; the interface refreshes the route and displays a warning when no route is available.

4.3. Fire-Spread Model

The fire-spread routine updates the map every 3 s: each fire cell converts its passable four-neighbor cells into fire zones, and original cell states are stored for restoration. Replanning starts after every update, so the route follows the current hazard field. The four-neighbor time-step representation captures directional and time-dependent growth on the navigation grid and is connected to established fire-dynamics concepts [9]. Section 5.3 evaluates 1, 3, 5, and 10 s update intervals over a 60 s scenario.

5. Experiment and Results Analysis

The evaluation covers planner comparison, shared-timeline ablation, fire-spread interval sensitivity, map-size scalability, and group-disjoint detector testing. Path cost is len(path) − 1 in four-neighbor grid steps; these entries are grid-step counts rather than physical distances. Planning time covers a single complete A* call unless a value is explicitly identified as cumulative. Planning success means that a valid start-to-exit route exists on the hazard map supplied to the planner. Fire-free and buffer-clear rates are evaluated against the shared actual-fire map: a fire-free path has zero overlap with fire, whereas a buffer-clear path remains outside the non-fire portion of the 3 × 3 buffer.

5.1. Comparison of Path-Planning Algorithms

A*, Dijkstra, D* Lite and RRT used the same 41 × 41 map, start-exit pair, hazard map and four-neighbor movement model. Each algorithm received 30 measured calls after five warm-up calls, with a fixed RRT seed per trial. Table 3 reports single-call planning time, path cost, and success rate; all runs used an Intel Core i5-14600KF CPU. These values refer to one fixed hand-designed map.

Table 3. Planner comparison on one fixed hand-designed 41 × 41 grid (30 measured calls after five warm-ups; mean ± sample SD).

Algorithm	Planning Time (ms, mean ± SD)	Path Cost (Grid Steps, mean ± SD)	Success Rate (%)
A*	0.407 ± 0.011	38.0 ± 0.0	100.0
Dijkstra	0.493 ± 0.026	38.0 ± 0.0	100.0
D* Lite	0.966 ± 0.051	38.0 ± 0.0	100.0
RRT	94.939 ± 48.451	62.6 ± 29.6	100.0

A*, Dijkstra, and D* Lite each returned the 38.0-step minimum-cost path in every trial with 100% success. A* required 0.407 ± 0.011 ms, Dijkstra 0.493 ± 0.026 ms, and D* Lite 0.966 ± 0.051 ms. RRT achieved 100% success with a path cost of 62.6 ± 29.6 steps and a planning time of 94.939 ± 48.451 ms. A* provided the lowest single-call planning time and the minimum observed path cost among the tested planners.

5.2. Shared-Timeline Ablation Experiment

One 60 s actual-fire sequence with updates every 3 s was reused for all ablation configurations, giving 20 common replanning events. Only the hazard information supplied to A* differs: the full system receives current fire and the 3 × 3 buffer, the without-buffer configuration receives current fire only, and the without-spread configuration receives only the initial fire map. Planner success is evaluated on the supplied map, whereas fire-free and buffer-clear rates are evaluated against the shared actual-fire map. All configurations therefore face identical hazard evolution, and the three rates are reported separately rather than treated as substitutes (Table 4).

Table 4. Shared-timeline ablation over 20 replanning events in the 60 s scenario (same actual-fire sequence in all configurations; fire-free and buffer-clear rates use the shared ground-truth map).

System Configuration	Planner Success (%)	Fire-Free Path Rate (%)	Buffer-Clear Path Rate (%)
Full system	30.0	30.0	30.0
Without 3 × 3 buffer	35.0	35.0	25.0
Without fire-spread information	100.0	15.0	10.0

The full system returned valid paths at 6 of 20 events (30.0%); all were fire-free and buffer-clear, with a mean path cost of 48.7 ± 8.3 steps and a mean minimum fire distance of 2.17 ± 0.41 cells. Without the buffer, planner success was 35.0%, buffer-clear rate was 25.0%, two paths entered the shared buffer, and the mean minimum distance decreased to 1.86 ± 0.69 cells. Without spread information, A* returned paths at all 20 stale-map events, while 15.0% were fire-free and 10.0% were buffer-clear; 17 paths intersected the shared fire, and 18 entered its buffer. The stale-map fire field differed from the shared actual fire by 76.3 cells per event on average (maximum 192), and the mean minimum fire distance was 0.30 ± 0.80 cells. These results show that current hazard information and the 3 × 3 buffer preserve the fire-clearance margin in the tested scenario.

5.3. Sensitivity Analysis of Fire-Spread Intervals

The sensitivity analysis used the same 41 × 41 map, start and stair-region goals over 60 s. Fire intervals of 1, 3, 5, and 10 s advanced the actual fire by one four-neighbor step at each event, rebuilt the 3 × 3 buffer, and triggered A* replanning from the fixed start. Path cost and planner timing follow the definitions used in Section 5.1.

As illustrated in Figure 3, the 1, 3, 5 and 10 s settings produced 60, 20, 12 and 6 scheduled events, final fire regions of 817, 194, 80 and 31 cells, and valid-path ratios of 10.0%, 30.0%, 50.0% and 100.0%. A valid path is free of both fire and the 3 × 3 buffer. Across successful events, the mean path cost was 48.7 ± 8.3 steps. Mean A* time increased from 0.038 ms per event at 1 s to 0.377 ms per event at 10 s, while cumulative A* time remained between 2.26 and 2.29 ms. Per-event and cumulative times are reported separately because the four intervals produced different numbers of events. The 3 s setting provided 20 hazard updates and a 30.0% valid-path ratio in the 60 s scenario.

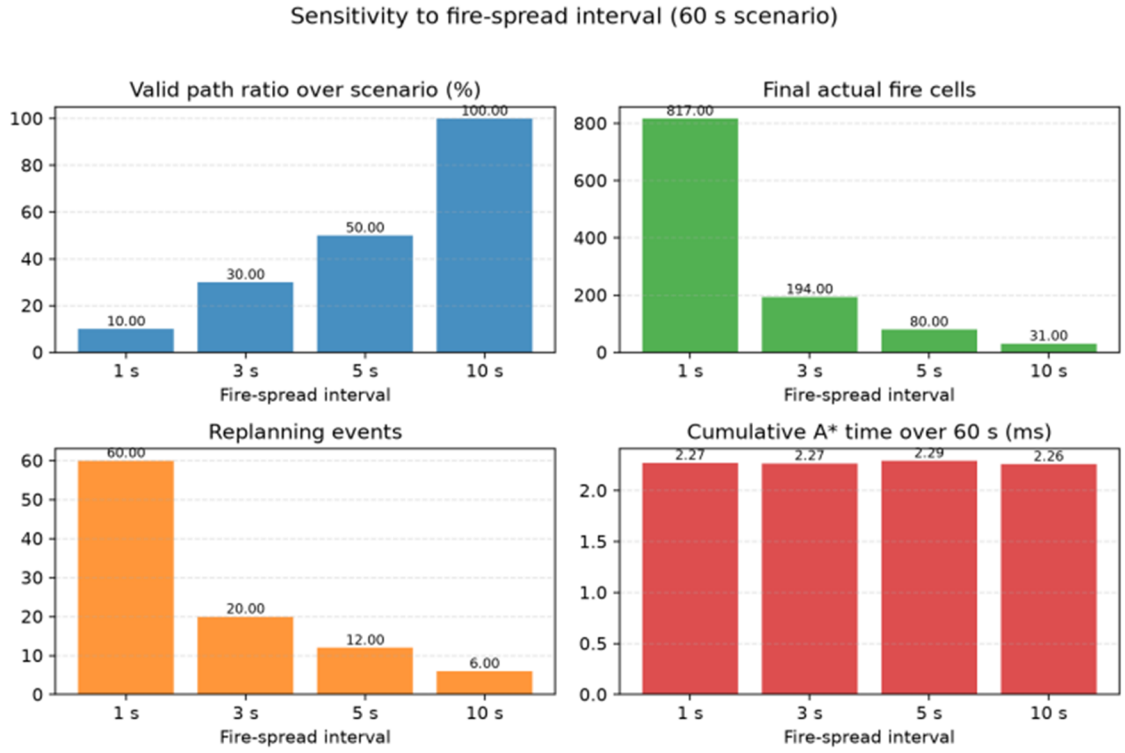


Figure 3. Sensitivity metrics under four fire-spread intervals in a 60 s scenario.

A common-timeline evacuation check used a 0.5 m cell length and a nominal walking speed of 1.2 m/s, giving 0.417 s per four-neighbor step. Occupant movement and fire updates used the same simulation clock, and a blocked remaining route triggered replanning from the occupied cell. For all four intervals, the occupant followed the 38-step (19.0 m) route and reached the stair region at 15.83 s; at the 3 s update interval, this arrival preceded the first loss of a valid path from the original start at 21 s. This check measures arrival relative to route availability and is separate from the valid-path ratio above.

5.4. Multi-Scale Map Test

For each map size of 21 × 21, 41 × 41, 61 × 61, and 81 × 81, 30 independently generated connected maps were evaluated with a start at the map center and a bottom-right exit region. Each map received three warm-up

runs and five timed A* calls; timing covers initialization, search, and path reconstruction. These generated maps and their center-to-exit setup differ from the fixed benchmark in Section 5.1.

All 120 generated maps produced valid four-neighbor paths, giving 100% planning success. Mean A* time increased from 0.133 ± 0.072 ms at 21 × 21 to 3.073 ± 1.707 ms at 81 × 81, and mean path cost increased from 26.6 ± 9.6 to 204.3 ± 95.9 steps. Table 5 uses generated maps with a center start and a bottom-right exit region, whereas Table 3 uses one fixed hand-designed map with a different start-exit definition; their 41 × 41 entries therefore represent distinct experimental scopes.

Table 5. Multi-scale A* results on 30 generated connected maps per size (mean ± SD; center start and bottom-right exit region).

Map Size	A* Planning Time (ms, mean ± SD)	Path Cost (Grid Steps, mean ± SD)	Success Rate (%)
21 × 21	0.133 ± 0.072	26.6 ± 9.6	100.0
41 × 41	0.667 ± 0.353	84.8 ± 36.6	100.0
61 × 61	1.475 ± 0.968	146.4 ± 90.2	100.0
81 × 81	3.073 ± 1.707	204.3 ± 95.9	100.0

5.5. Group-Disjoint Independent Detector Test

YOLO11m was evaluated on 497 original real-world flame images from 327 unseen groups, containing 899 annotated fire boxes. Using a 512 × 512 input, a confidence threshold of 0.001, IoU 0.7, and a maximum of 300 detections, the model achieved 69.46% precision, 64.50% recall, 72.36% mAP@0.5 and 46.47% mAP@0.5:0.95. Precision and recall are reported at the maximum-F1 operating point, whereas the mAP values are obtained by integrating the full precision–recall curve. The 72.36% result is the final group-disjoint out-of-sample estimate and is not pooled with the validation or architecture-comparison values in Tables 1 and 2.

6. Conclusions

This paper presented an integrated fire escape guidance system that combines YOLO11 flame detection, map-aligned hazard grids, and event-triggered A* replanning. Detected flames are mapped to grid cells, expanded by a 3 × 3 safety buffer, updated through time-dependent fire growth, and converted into evacuation routes and alarms. YOLO11m achieved 72.36% mAP@0.5 on the group-disjoint independent test set, and A* averaged 0.407 ± 0.011 ms per call. Shared-timeline, interval-sensitivity, map-scale, and independent real-world image experiments validate the integrated workflow. The contribution is the detector-to-grid interface, the event-triggered closed-loop replanning workflow, and the system-level validation of this integration.

Author Contributions

J.A.: Conceptualization, methodology, writing—original draft, writing—review & editing. L.G.: Methodology, software, formal analysis. X.L.: Methodology, software, formal analysis. L.Z. and Q.F.: Methodology, software, formal analysis, data curation, investigation. T.S.: Software, validation, investigation. All authors have read and agreed to the published version of the manuscript.

Funding

The study is supported by the National Natural Science Foundation of China with funding number 62506243 and the Youth Science Foundation of Liaoning Provincial Department of Education with project number LJ212510144013.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

The curated flame-image dataset, the 302-frame flame recording, and the generated grid maps are available from the corresponding author upon reasonable request for non-commercial academic research, subject to the terms and permissions of the original data providers.

Conflicts of Interest

The author Q.F. was employed by NARI Technology Company Limited. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Use of AI and AI-Assisted Technologies

No AI tools were utilized for this paper.

References

1. Danish, S.; Piran, M.J.; Khan, S.U.; et al. Vision-based fire management system using autonomous unmanned aerial vehicles: A comprehensive survey. *Artif. Intell. Rev.* **2026**, *59*, 16. https://doi.org/10.1007/s10462-025-11415-3.
2. Tao, H.; Huang, Z.; Wang, Y.; et al. Efficient feature fusion network for small objects detection of traffic signs based on cross-dimensional and dual-domain information. *Meas. Sci. Technol.* **2025**, *36*, 035004. https://doi.org/10.1088/1361-6501/adb2ad.
3. Gao, L.; Zhuang, Z.; Tao, H.; et al. Non-lifted norm optimal iterative learning control for networked dynamical systems: A computationally efficient approach. *J. Frankl. Inst.* **2024**, *361*, 107112. https://doi.org/10.1016/j.jfranklin.2024.107112.
4. Shirwaikar, R.; Narvekar, A.; Hosamani, A.; et al. Real-time semi-occluded fire detection and evacuation route generation: Leveraging instance segmentation for damage estimation. *Fire Saf. J.* **2025**, *152*, 104338. https://doi.org/10.1016/j.firesaf.2025.104338.
5. Hart, P.; Nilsson, N.; 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.
6. Wu, S.; Chen, H.; Zhang, Y.; et al. Robot visual servo based on lightweight YOLO11-SMMA for Camellia oleifera fruits harvesting. *Comput. Electron. Agric.* **2026**, *243*, 111409. https://doi.org/10.1016/j.compag.2026.111409.
7. Fang, H.; Zhang, B.; Tam, W.C.; et al. A deep learning-based approach for unsafe area prediction in building fire evacuation. *J. Build. Eng.* **2025**, *112*, 113723. https://doi.org/10.1016/j.jobe.2025.113723.
8. Fang, H.; Tam, W.C.; Lovreglio, R.; et al. Development of a tenability-based path planning model for building fire evacuations using reinforcement learning. *J. Build. Eng.* **2026**, *118*, 115132. https://doi.org/10.1016/j.jobe.2025.115132.
9. Drysdale, D. *An Introduction to Fire Dynamics*, 3rd ed.; John Wiley & Sons: Chichester, UK, 2011. https://doi.org/10.1002/9781119975465.