applied sciences

MDPI

Article

A Hybrid Sensor-Fusion and Graph-Based Routing Model for Explainable Multi-Agent Fire Evacuation

Rosen Ivanov

Department Computer Systems and Technologies, Technical University of Gabrovo, 5300 Gabrovo, Bulgaria; rosen@tugab.bg

Abstract

Fire evacuation guidance in multi-storey buildings must adapt to spreading fire and smoke, but adaptive machine-learning methods are hard to interpret and verify in regulated safety applications. This paper presents a deterministic alternative that extends the hierarchical building-graph model with sensor fusion with hysteresis, hard exclusion of hazardous nodes, risk-first route ranking, coordination of agents through virtual load, and a cellular-automaton movement model. Only one part of the model is fitted to data: a linear cost model with eight coefficients. It is used only to break ties between routes. All other parameters are set either by regulation or by engineering judgement. As a result, every recommendation can be traced back to sensor values and formulas. Thirteen simulated scenarios in four buildings, with up to 120 agents, were evaluated against four ablated variants with 20 Monte Carlo replicates each. Hard exclusion ensures that recommended routes contain no node classified as fire- or smoke-affected at planning time, which held in all scenarios for agents with an admissible route; this is a model-level property that depends on sensor classification and does not cover agents who do not evacuate. With a fixed coordination coefficient ($\lambda = 0.35$), coordination increased the number of exits used in eight scenarios, at a cost in maximum evacuation time in most scenarios where the difference was significant. In the two largest scenarios (120 agents), the standard deviation of the maximum evacuation time was 2.7 to 4.9 times lower than for a baseline with neither coordination nor calibrated weights, but not lower than for the variant without coordination, and the full system had the lowest variance in eight of thirteen scenarios. The only contacts with fire or smoke during movement (Scenario 12) resulted from sensor-classification latency. Code and data are provided for full reproduction.

Keywords: fire evacuation; deterministic routing; sensor fusion; risk assessment; regulatory compliance; multi-agent systems; dynamic route planning

Check for updates

Academic Editor: Paolo Renna

Received: 17 October 2026
Revised: 29 October 2026
Accepted: 2 November 2026
Published: 9 November 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Fires in multi-storey and architecturally complex buildings continue to represent one of the most serious challenges for human-safety systems. Regardless of advances in building codes, fire-protection installations, and early-warning systems, effective evacuation remains a critical factor in limiting human losses. Our previous work [1] proposed a hierarchical building-graph model for evacuation guidance, which this paper extends with a deterministic, traceable routing approach. Data from the systematic bibliometric analysis by Qi et al. [2], covering 673 publications from 2000 to 2024, show that research interest in fire evacuation is growing exponentially. 59.3% of all publications are concentrated in the last five years, reflecting the increasing recognition of the problem as an unresolved

scientific task. The main reason is that real fires are highly dynamic events in which building conditions change continuously because of fire spread, smoke accumulation, decreasing visibility, changes in air toxicity, and the emergence of local congestion along evacuation routes. Under such conditions, predefined evacuation plans often cannot guarantee optimal, or even safe, behavior for all occupants.

Over the last decade, a significant transformation has been observed in fire-evacuation research. According to the systematic reviews by Muthukumar et al. [3] and Qi et al. [2], the field is gradually moving from static evacuation models toward intelligent systems that combine sensor data, digital building models, fire-development simulations, and artificial-intelligence algorithms. Muthukumar et al. [3] analyze 96 peer-reviewed publications and identify a clear shift from rule-based systems to deep-learning solutions, noting that the largest share of AI research (43.75%) is directed at the disaster-response phase—the phase into which real-time evacuation routing falls. This trend is driven by the need for evacuation decisions to be adapted to the specific real-time situation rather than based on predefined scenarios.

In parallel, technologies for creating digital representations of the building environment are also developing. The integration of Building Information Modeling (BIM) and Fire Dynamics Simulator (FDS) enables the construction of digital twins that combine the building's spatial structure with fire-dynamics models and risk assessment [4]. Simulation studies by Zhang et al. [5] in L-shaped buildings demonstrate that the safety of evacuation routes depends not only on distance to exits but also on smoke dynamics, airflow direction, and fire-origin location—factors whose consideration requires a substantially richer information basis than that provided by static plans.

Significant progress is also observed in the use of machine-learning and reinforcement-learning algorithms. Contemporary research demonstrates the potential of deep learning for predicting fire development, assessing risk, and dynamically planning evacuation routes [6–8]. In addition, the development of multi-agent approaches enables coordination of the behavior of large groups of people and management of evacuation processes in complex environments [9–11].

Despite these significant achievements, there is a substantial contradiction between the adaptivity of contemporary AI-based solutions and the requirements for reliability, transparency, and traceability in safety-critical systems. Most contemporary models achieve high effectiveness through complex neural architectures whose behavior is difficult to interpret and formally verify. From the standpoint of real-world deployment, this represents a serious limitation, since evacuation decisions in safety-critical systems must not only be effective but must also meet the requirements of the applicable regulations and fire-safety building codes, which presuppose deterministic, traceable, and verifiable system behavior. Muthukumar et al. [3] and Qi et al. [2] identify precisely the lack of explainability, limited generalizability, and difficulties in real-world deployment as among the field's main unresolved problems. These limitations become especially critical when a system must be approved by regulatory bodies or certified for real-world operation.

The present study treats evacuation as a real-time intelligent decision-support task in which adaptivity to a dynamically changing environment is achieved without the use of trainable or "black-box" components, relying entirely on deterministic, configurable, and traceable mechanisms. This work is developed as a direct extension of the BFS-based KSP algorithm for hierarchical buildings proposed by Ivanov [1], whose deterministic graph architecture is retained as the system's core and extended with three complementary, fully formula-based mechanisms: a calibrated model for weighting graph edges, threshold-based adaptation of priorities, and multi-agent coordination.

The main contributions of the present work are:

- A fully deterministic evacuation algorithm that combines BFS/K-shortest-paths search with an externally configurable and fully auditable composite edge-cost model, a threshold-based table for adaptive node priority, and multi-agent coordination. Sections 4.4 and 4.5 show empirically that two things do not depend on the exact numerical calibration of the edge-cost model: the exclusion of classified fire and smoke nodes from recommended routes, and the measured adaptivity. This is a deliberate architectural property, not a limitation, because it means the safety property does not rely on precise numerical tuning. The choice among equally risky routes is different: it still depends on the secondary cost term, which Section 4.4.5 discusses;
- A weighted sensor-fusion module with hysteresis for converting continuous sensor readings (temperature, smoke, CO_2, crowd density) into topological-graph parameters, without trainable components;
- A cellular automaton with directed diffusion and a danger-escape mechanism for realistic modeling of evacuees under a dynamically spreading fire;
- An architectural separation between statically calibrated parameters and the dynamically adapted priority layer, providing both full traceability/auditability and adaptivity, a prerequisite for regulatory compliance and certification in real evacuation system
- Empirical validation across thirteen scenarios in four real buildings, including ablation configurations, one-at-a-time sensitivity analysis, and a Monte Carlo comparative study against multiple baseline configurations and alternative route-sorting criteria.

The remainder of the paper is organized as follows: Section 2 reviews the existing literature and positions the proposed approach relative to it; Section 3 describes the methodology—the architecture, risk assessment, the route-cost model, adaptive node priority, agent coordination, and the crowd model; Section 4 presents comparative experiments across thirteen scenarios in four buildings; Section 5 verifies compliance with regulatory requirements; Section 6 discusses the results and limitations; Section 7 summarizes the contributions and conclusions.

2. Related Work

To situate the present study within contemporary trends, it is necessary to review the available systematic reviews that outline the field's development and identify its unresolved problems.

Muthukumar et al. [3] present a systematic review of AI applications in disaster management, analyzing 96 peer-reviewed publications from IEEE Xplore, Elsevier, Wiley, and Springer using the PRISMA methodology and bibliometric analysis with VOSviewer (1.6.20). The review establishes a clear evolution from rule-based systems to machine-learning and deep-learning solutions. The authors identify three key unresolved problems: (1) the quality and representativeness of training data; (2) the limited generalizability of models to unfamiliar scenarios and environments; and (3) the substantial gap between laboratory developments and real-world operational deployment in emergency systems. These problems directly motivate architectural decisions in the present work.

Qi et al. [2] conduct a bibliometric and systematic analysis of 673 publications in the field of fire evacuation for the period 2000–2024 and identify three dominant research directions: crowd-behavior analysis, evacuation simulation and optimization, and risk perception in decision-making. The authors propose the "perception–behavior–model" framework as a unifying conceptual basis and identify the growing use of AI and RL for dynamic optimization after 2021 as a leading trend. A persistent problem is the shortage of real fire data, which forces most studies to rely on simulated data.

The two reviews outline the boundaries of current knowledge and serve as a methodological basis for structuring the remaining subsections of this review.

2.1. Modeling the Fire Environment

The effectiveness of any evacuation-support system depends directly on the quality of information about building conditions and fire development. For this reason, a significant portion of contemporary research is directed toward building digital representations of the building environment that enable risk assessment and prediction of incident development.

Zhang et al. [5] use the Fire Dynamics Simulator (FDS) to analyze nine fire scenarios in a 15-storey L-shaped building, systematically varying fire origin, floor, and wind direction. The results demonstrate that visibility is the most sensitive indicator of fire development. It responds faster than temperature and CO and shows a larger hazardous zone. More importantly, the authors find that the safety of evacuation routes depends nonlinearly on the combination of all these factors: in the baseline scenario, Stairwell 1 remains safe for only about 60 s, whereas under different wind conditions the order of preference is completely reversed. This shows that traditional static evacuation plans are structurally unable to guarantee optimal decisions during real, evolving fires.

Leong et al. [4] extend the capabilities of FDS through integration with Building Information Modeling (BIM). Their proposed conceptual framework, implemented on a three-storey building in Revit, the Dynamo API, and PyroSim, enables automated regulatory-compliance checking and the integration of fire-simulation results back into the BIM model. From the simulation, an Available Safe Egress Time (ASET) of 331 s was determined, while the evacuation time was under 30 s. The work demonstrates a functioning pipeline from BIM through fire scenario to evacuation decision. In this way, BIM becomes not merely a design tool but also a platform for risk analysis and safety management.

Further development of this idea is presented by Deng et al. [12], who use an ontological model (IFEEont) to represent fire-safety knowledge, integrated with BIM through IFC and Protégé/SWRL rules. The approach enables automated semantic inference of space-accessibility information and dynamic generation of evacuation routes, validated on a six-storey office building in Guangzhou. The method improves interoperability between different information systems but offers no trainable components and relies on predefined rules.

In the same direction, Jiang et al. [13] propose a method for semantically enriching a BIM model, specifically targeting the requirements of fire simulation. Their method automatically populates the BIM model with information needed by the simulation engine (material properties, boundary conditions, space classification) that is typically missing from design-oriented BIM models. The approach demonstrates optimization of the evacuation design of public buildings directly on the enriched model, confirming that the quality of the BIM-to-simulation conversion remains a key bottleneck regardless of the chosen evacuation algorithm itself.

2.2. Dynamic Evacuation Route Planning

Classical studies in this area treat evacuation as the task of finding an optimal route between a starting point and a destination in the presence of dynamically changing constraints.

Shih and Tsai [14] propose a multi-criteria approach for determining the safest fire-evacuation route, in which fire-spread prediction via a Context-Embedding and Multi-Scale Spatiotemporal-Expression LSTM (CMS-LSTM) model is combined with the SPS (Safest Path Search) algorithm. The system achieves over 10% higher survivability compared with previous methods and is about 10 times faster at computing routes, thanks to a two-level map architecture. The limitation is that effectiveness depends directly on the accuracy of the fire prediction.

Guan et al. [15] develop this idea further through a bi-objective optimization model that simultaneously minimizes risk and path length via a modified Dijkstra algorithm.

Risk is assessed through five factors (CO, HCN, temperature, visibility, crowd density), combined through linear weighting. The results show that as the fire develops, the system successfully reroutes from a shorter, riskier path to a longer but safe one.

Yen et al. [16] develop two algorithms: Smoke-aware Individual Evacuation Path (SIEP) for individual evacuation and Smoke-aware Group Evacuation Path (SGEP) for group evacuation in an IoT-equipped multi-storey building. The study is particularly relevant to the present work along three axes: (1) IoT sensor integration for real-time monitoring of smoke toxicity, analogous to the proposed sensor-fusion module; (2) the Fractional Effective Dose (FED) indicator as a quantitative measure of route risk, paralleling the proposed hazard function; and (3) prioritized group evacuation based on cumulative toxic exposure, analogous to the proposed multi-agent coordination. SGEP outperforms SIEP by up to 68.1 percentage points in a methane fire and by 50% in a polyvinyl chloride (PVC) fire. The key difference from the proposed approach is that Yen et al. use binary node classification without continuously weighted, calibrated risk scores and without hysteresis, which makes the system more vulnerable to sensor noise.

Li et al. [17] focus on the interaction between fire spread and evacuee behavior in multi-storey hybrid buildings. The study is methodologically distinctive in that it uses Stacking Ensemble Learning (XGBoost + LightGBM + Gradient Boosting with a Lasso meta-model) to predict evacuation time, trained on 4520 simulation records from Pathfinder. The model achieves Mean Absolute Error (MAE) = 88.85 s and Mean Relative Error (MRE) = 18.88% on the test data, while the integrated routing model reduces total evacuation time by 7.12% for 4520 evacuees. This paper is closer to Section 2.3, since it uses machine learning, and highlights the potential of predictive methods to support route selection.

Bai et al. [18] propose an improved A* algorithm with dynamic fire-risk assessment, integrating temperature, visibility, CO, and O_2 into a single quantitative score. Validation in a metro station shows about a 69–70% reduction in route risk compared with standard A*, at a moderate increase in path length (12–46%). The approach is methodologically close to the proposed algorithm in terms of its dynamic risk function, but it does not include sensor fusion with hysteresis, agent coordination, or crowd modeling.

Ant-colony-based algorithms also occupy an important place. Xu et al. [19] develop an improved Ant Colony Optimization (IACO) for a single-storey supermarket, in which the equivalent path length incorporates temperature, CO, and visibility. Zhang et al. [20] propose adaptive variants of IACO (IAACO) with a multi-factor heuristic function (risk, energy cost, length) and an adaptive evaporation coefficient, achieving up to a 63.8% reduction in inefficiency relative to standard ACO in environments with multiple exits and fire origins. Both approaches demonstrate a good ability to reach a global optimum but require substantial computational resources when scaled.

Zuo et al. [21] use Dynamic Avoid-Smoke A-Star (DASA-Star) for evacuation in metro stations, achieving a 66.8% reduction in fire risk relative to Dynamic A-Star at only a 9.2% increase in route length. Golshani and Fang [22] develop a navigation model integrating FDS, modified Dijkstra, Agent-Based Simulation, and Intelligent Adaptive Dynamic Signage (IADSS), accounting simultaneously for travel time, fire impact (FED), and congestion. Theirs is one of the few publications to consider all three competing criteria in parallel.

Regardless of the differences among these approaches, their common feature is that evacuation is treated predominantly as an optimization task with predefined rules. In most cases, a mechanism for continuous learning and adaptation to unfamiliar scenarios is absent.

2.3. Reinforcement Learning for Evacuation

As computational capabilities grow, increasing attention is being paid to reinforcement learning as a means of decision-making in a dynamic environment.

Sharma et al. [6] make two independent contributions: (1) they develop the first standardized RL environment for fire evacuation in the form of an OpenAI Gym with dynamic spread, uncertainty, and capacity constraints; and (2) they propose Q-matrix transfer learning, in which tabular Q-learning pretrains a DQN network. The results show significant improvement. QMP-Dueling DQN achieves an average of 68.75 steps compared with 642 for SARSA and 343 for PPO. An important methodological contribution is precisely the division of the task into a "structural-knowledge transfer" phase and an "online adaptation" phase. This principle is analogous to the architectural separation in the present work between the statically calibrated route-cost model and the dynamically adapted node-priority table, albeit without a trainable component on the side of the proposed system.

Feng et al. [7] use an improved DDPG algorithm (I-DDPG with RND, HER, and BWO) for continuous control of evacuation decisions in a hospital building, achieving about 100 units higher average reward and faster convergence relative to standard DDPG. Unlike Sharma et al., this work is directed at algorithm optimization rather than at a new environment. Its limitation is that the model considers a single agent without interactions among people.

Fang et al. [8] introduce the concept of tenability-based path planning through the SafeStep model, in which the reward function is based on the FED of toxic gases rather than solely on geometric characteristics. In a test scenario, SafeStep achieves a 62% reduction in FED exposure relative to Dijkstra's algorithm and successfully identifies evacuation routes in scenarios in which Dijkstra fails to find a solution. Cumulative accounting of toxic exposure is an important step toward more realistic risk modeling.

Despite the high adaptivity of reinforcement learning, these approaches suffer from limited interpretability. Decisions are formed by complex neural models, which makes it difficult to explain the chosen route and raises questions about their reliability in safety-critical applications.

2.4. Multi-Agent Systems and Coordinated Evacuation

Under real conditions, evacuation is rarely reducible to the behavior of a single individual. For this reason, recent years have seen growing interest in multi-agent models.

Sun et al. [9] propose GRU-MAPPO, an architecture combining Multi-Agent Proximal Policy Optimization with a Gated Recurrent Unit for processing temporal dependencies. Agents are organized hierarchically (leaders and followers), with leaders planning global routes via GRU-MAPPO while local collision avoidance is handled through Optimal Reciprocal Collision Avoidance (ORCA). The method achieves about 20% faster convergence relative to standard MAPPO, a Path Efficiency of 0.91, and a Uniformity of 0.91 when tested on a real building floor plan.

Xie et al. [10] extend this concept by integrating fire-spread models and dynamic evacuation signage within the Cell Transmission Model (CTM)–PyroSim–QMIX framework. Each dynamic sign acts as an agent coordinated by the QMIX algorithm. Across 10 fire scenarios in a real karaoke-complex floor plan, the method achieves maximum improvements of over 30% in evacuation efficiency, over 50% reduction in fire risk, and over 70% reduction in congestion risk relative to static signage.

Omidbeyk and Ghavami [11] combine GIS, BIM, and Proximal Policy Optimization (PPO)-based reinforcement learning within a unified multi-agent simulation environment implemented in Unity ML-Agents. The study demonstrates that fire location is a determining factor. For example, for a fire on the sixth floor versus a basement level, evacuation time

increases by 57%. RL-trained agents consistently outperform those following a static shortest-path strategy. Spatial context from BIM/GIS has a substantial effect on the quality of the decisions made.

Shi et al. [23] specifically address the problem of balanced use of multiple exits through a Collaborative Decision-Making Framework combining Multi-Agent Proximal Policy Optimization with a dynamic, multi-objective reward mechanism and a localized (rather than global) observation strategy for each agent. The authors achieve over 70% more stable training and a 2–11% reduction in mean evacuation time relative to baseline models in an asymmetric multi-exit scenario.

Beyond purely algorithmic improvements, Yu et al. [24] demonstrate an application of multi-agent modeling to a real outpatient hospital building, aimed at resource-allocation decisions for fire response and prioritization of retrofit measures. They show that multi-agent models can serve not only for algorithm evaluation but also for direct support of engineering retrofit decisions. Bahmani et al. [25] extend the scope of multi-agent modeling to a specific, vulnerable occupant group (children in elementary schools), incorporating age-specific psychomotor and health characteristics into the agents.

These approaches demonstrate high adaptivity and an expanding range of applications—from algorithmic optimization to specific engineering and demographic contexts—but they increase model complexity and further complicate the analysis of the reasons behind specific decisions.

2.5. Integrated AI Platforms for Disaster Management

The most recent trends are directed toward building integrated systems that combine multiple AI technologies.

Saravanan et al. [26] propose architecture including computer vision (OpenCV + Canny Edge Detection), a Social Force Model for crowd simulation, Retrieval-Augmented Generation (RAG) with FAISS and GPT-2 for generating context-dependent evacuation policies, and a Graph Attention Network (GAT) for predicting optimal routes. The system achieves an average route efficiency of 0.93 with zero incursions into the hazardous zone around the fire. The limitation is that validation is entirely on synthetic scenarios and publicly available floor plans.

In a direction opposite to computational complexity, Yang et al. [27] investigate the integration of a large language model (LLM) as the "cognitive core" of each agent in an agent-based model (ABM) simulation of disaster evacuation. Each agent maintains personality traits, observations, and a decision history, while the LLM provides context-aware reasoning and inter-agent communication. The improvements demonstrated in realism and adaptivity come at the cost of complete opacity in the agent-level decision-making process.

A systematic literature review by Lin et al. [28], covering the period 2004–2025 using the PRISMA methodology, categorizes integrated digital-twin-type evacuation platforms into three functional components: sensing and environment reconstruction, disaster simulation, and planning/risk assessment. This is a framework within which the architecture of the present work is positioned precisely in the third component, but without dependence on an opaque digital twin. In addition, Pishahang et al. [29] propose a Bayesian agent-based model is implemented as a complete software tool for wildfire evacuation planning. This demonstrates that integrating probabilistic uncertainty with directly applicable software is a distinct direction of development, separate from the AI-centric platforms discussed above.

Sun et al. [30] directly address the problem of decision-making time through EvacuNet, a deep neural network trained on 5400 simulated evacuation scenarios spanning diverse building types, exit configurations, and density distributions. EvacuNet can provide an estimate of evacuation time and exit-wise distribution in under 5 s with over 90% accuracy,

outperforming even human experts in speed, consistency, and accuracy. Unlike conventional simulations, which require hours, or other AI-based models, which require minutes, EvacuNet pushes the boundary of practical applicability of such tools toward genuine real time.

These examples, spanning computer vision and RAG, LLM-based reasoning, systematized digital-twin frameworks and Bayesian software tools, and deep-learning models for real-time assessment, show a trend toward increasingly powerful, faster, and methodologically more diverse, yet increasingly harder-to-audit architectures. This trend stands in direct contrast to the goals of the present work.

2.6. Positioning and Scientific Novelty of the Proposed Approach

The literature review conducted outlines three main directions. Classical graph-based and optimization algorithms [14–16,18–22] provide high traceability and determinism but rely on statically defined rules and adapt poorly to unforeseen scenarios. RL- and MARL-based solutions [6–11] achieve high adaptivity, but their decisions stem from trained policies that are difficult to interpret, which complicates auditing and practical certification. Integrated BIM/IoT/AI platforms [4,5,12,26] improve the information basis but do not resolve the problem of decision transparency.

This gap is especially critical for systems with direct relevance to human life, where applicable regulatory requirements and fire-safety standards presuppose the ability to fully verify and justify every generated decision.

A summary comparison of the proposed approach against these three directions, in terms of dynamism, logic type, explainability, graph structure, real-time capability, multi-agent coordination, and potential for regulatory compliance, is presented in Table 1.

Table 1. Comparison of the proposed approach with representative methods from literature.

Approach	Dynamic Evacuation	Logic Type	Explainability	Graph Structure	Real Time	Multi-Agent Coordination	Potential for Regulatory Compliance
Classical routing algorithms [14,15]	Partial	Deterministic (rule-based)	High	Yes	Limited	No	High
ACO/A*/risk-oriented approaches [16,18–21]	Yes	Deterministic (rule-based)	Medium	Yes	Yes	No	Medium
Single-agent Reinforcement Learning [6–8]	Yes	Trainable	Low	No	Yes	No	Low
Multi-Agent Reinforcement Learning [9–11]	Yes	Trainable	Low	Partial	Yes	Yes	Low
BIM/FDS [4,5,12]	Partial	Deterministic (rule-based)	High	Yes	Limited	No	Medium
Integrated AI platforms [26]	Yes	Trainable	Low	Yes	Yes	Partial	Low
Proposed approach	Yes	Deterministic	High	Yes	Yes	Yes	High

Of all the approaches we reviewed, ours is, to our knowledge, the only one that combines dynamic adaptivity and multi-agent coordination with fully deterministic, calibrated logic. Table 1 is a qualitative comparison, and it is not based on a common benchmark.

At the component level, the novelty consists of: a continuous risk-assessment model based on weighted fusion of heterogeneous sensor data with hysteresis (instead of predefined fixed binary thresholds, as in Yen et al. [16], or static rules, as in Bai et al. [18]); a calibrated composite model for graph-edge costs, whose coefficients are determined offline

through numerical calibration against reference scenarios rather than through statistical training; a threshold-based table for adaptive node priority, acting as a deterministic memory of accumulated evacuation experience; accounting for the projected future load through the cellular automaton when evaluating routes; and coordination among agents through dynamic virtual loading of already-occupied routes, rather than through trained policies (in contrast to Xie et al. [10]; Omidbeyk and Ghavami [11]; Sun et al. [9]).

At the architectural level, the main contribution is the development of a hybrid structure in which adaptivity to a dynamically changing environment is achieved without a single trainable component: the statically calibrated cost-model parameters and the dynamically adapted node-priority table work together, while the final routing remains entirely the outcome of deterministic graph search. This architectural separation between calibrated environment assessment and deterministic decision-making is the direct reason the system retains full traceability and reproducibility for every generated decision, including during post-incident audits.

At the system level, the novelty stems from the integration of a topological building model [1], real-time multi-sensor risk assessment, a calibrated route-cost model, threshold-based priority adaptation, deterministic routing, multi-agent coordination, and pedestrian-flow modeling. All of this is combined within a single, fully auditable platform designed to meet regulatory requirements for safety-critical systems.

For systems of this type, the difference in decision traceability relative to RL- and MARL-based approaches [6,7,9,10] is especially significant: whereas in those approaches the final route stems from a trained policy whose behavior cannot be directly justified by specific measurements, in the proposed approach every evacuation recommendation can be traced back to specific sensor values, risk-assessment formulas, and graph-model parameters—a property of direct practical significance for subsequent analysis, auditing, and regulatory approval.

3. Methods

3.1. General Overview of the Algorithm

The proposed system implements a hybrid architecture for supporting fire evacuation, combining real-time sensor-data processing, calibrated risk-assessment mechanisms, and deterministic graph-based routing. The main design goal is to simultaneously achieve adaptivity to a dynamically changing environment and preserve full traceability of the decision-making process. The algorithms are implemented in Python (3.9+) and are available for analysis and use (Supplementary Material/software). In addition to the source code, a supplementary tool has also been developed for interactive 3D visualization of a selected individual evacuation route (Section 5.6), useful for auditing a specific evacuation decision. The project structure, a description of the software, and the operation of the 3D visualizer are described in PDF files (Supplementary Material/docs).

The overall algorithmic architecture of the system is presented in Figure 1.

The architecture is organized as a sequence of interconnected modules that transform data from the building environment into concrete evacuation routes. The starting point of processing is the building's topological model, built on a BIM-based graph representation (IFC file format).

Data from the sensor infrastructure feeds into this model, including temperature, smoke, carbon-dioxide concentration, and evacuation-flow density. The sensor-fusion module converts the heterogeneous measurements into unified estimates of fire risk and smoke risk. The resulting values are used by the calibrated risk-assessment model, which determines the dynamic traversal costs for individual nodes and edges of the graph. In

parallel, the adaptive node-priority mechanism updates local environment assessments whenever new hazards arise or the fire configuration changes.

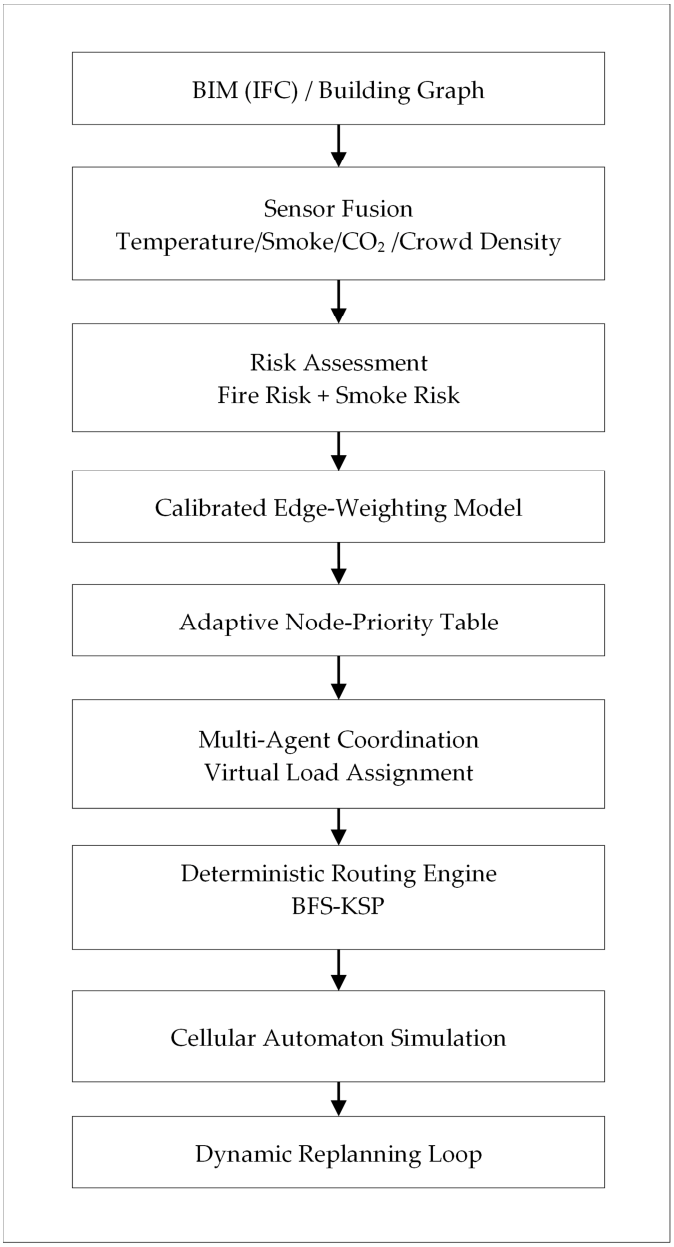


Figure 1. Algorithmic architecture of the proposed system.

Following risk assessment, the multi-agent coordination module runs, distributing evacuees among alternative routes through a virtual-load mechanism. Final route generation is performed by the deterministic BFS-KSP core, which uses the updated graph weights and environmental constraints to determine the most suitable evacuation paths.

The resulting routes are fed to the cellular automaton for modeling the evacuation flow, which simulates the movement of individual agents, their interactions, and the influence of local hazards on their behavior. The simulation results are used to update the load and risk estimates, closing the dynamic-replanning loop and enabling adaptation to the evolving fire situation.

The proposed algorithm implements multi-agent dynamic evacuation, in which agent routes are recomputed in real time whenever a significant change in hazard conditions

occurs, rather than being determined once at the start of the simulation. The algorithm operates in an iterative loop with a fixed five-second time step (the sampling period of the sensor series; sensor ingestion, hazard classification and the replanning trigger are evaluated at the same instant, and the value is a design choice, not derived from a standard or from measurements), at each step of which the following operations are performed in sequence: (1) incoming sensor data is normalized and combined into an updated hazard-state estimate for each node of the building graph; (2) this estimate is compared with the previous one to determine whether a significant change has occurred that requires route recomputation; (3) if a significant change is detected, K alternative routes are computed for each affected agent through weighted search over the connectivity graph, accounting simultaneously for risk, load, and distance; (4) the resulting route is passed to the crowd-movement model, implemented as a cellular automaton; and (5) agent positions and local crowd density are updated and recorded for the next step.

The remainder of this section describes in detail the mathematical apparatus behind each of these operations.

3.2. Hazard-State Assessment

Raw physical measurements (temperature, smoke level, CO_2 concentration, headcount) are transformed into normalized risk scores in the range [0, 1]. This is achieved through linear scaling relative to maximum values specified in a threshold configuration file (400 °C for temperature, 3000 ppm for CO_2; the smoke level arrives already normalized; the maximum crowding capacity depends on node type, as summarized in Table 2). These normalization maxima are engineering-judgement configuration values, not values derived from a cited standard. The CO_2 maximum is set well below toxicological danger thresholds, since it functions as a sensitive early-saturation point for the sensor signal—CO_2 rises rapidly with combustion long before reaching physiologically dangerous concentrations—rather than as a life-safety limit. The temperature normalization value (400 °C) is broadly consistent with the sub-flashover ceiling-jet temperatures obtainable from the fire-dynamics correlation used later in this study to generate the temperature time series (Section 4.4), though it was not formally derived from it.

Table 2. Maximum crowding capacity by node type, used for density normalization.

Node Type	Maximum Capacity (People)
Corridor	20
Stairwell	10
Room	15
Elevator	8
Safe zone (exit)	100
Undefined type (default)	15

The fire and smoke risk scores for each node are computed through weighted maximum aggregation of the normalized sensor values. Instead of simple summation, the maximum of the contributions from the individual sensor types is taken, weighted by their reliability for the corresponding hazard type (Equation (1)):

$$fire_risk(node) = \max\left(w_{f,T} \cdot norm_T,\ w_{f,S} \cdot norm_S,\ w_{f,CO_2} \cdot norm_{CO_2}\right), \quad (1)$$

where $fire_risk \in [0, 1]$ is the node's composite fire-risk score; $w_{f,T}$, $w_{f,S}$, w_{f,CO_2} are the confidence weights of the temperature, smoke, and CO_2 sensors for *fire* risk; $norm_T$, $norm_S$, $norm_{CO_2} \in [0, 1]$ are the corresponding normalized sensor readings (linearly scaled against

400 °C, 3000 ppm, and the smoke sensor's native $[0, 1]$ output, respectively, clipped to 1); and $\max(\cdot)$ takes the single most alarming weighted contribution rather than summing them.

An analogous formula, with a separate set of weights, is applied for $smoke_risk(node)$. The confidence weights for the individual sensor types wf are summarized in Table 3.

Table 3. Sensor-type confidence weights are used in the aggregation of fire risk and smoke risk.

Sensor Type	Fire Weight	Smoke Weight
Temperature	0.95	0.70
Smoke	0.70	0.85
CO_2	0.40	0.80
Crowd	0.00	0.00

For example, the temperature-sensor reading carries the highest weight for fire risk (0.95), while the smoke index carries the highest weight for smoke risk (0.85); crowd does not participate in this aggregation (weight 0 for both risk types), since crowding is treated as a separate, independent risk factor (see Section 3.4).

To avoid unstable switching of the classification for values near the threshold, the transition between a node's "safe" and "hazardous" state is performed via hysteresis: a node enters a hazardous zone at a higher threshold than is required to leave it again. The thresholds applied are summarized in Table 4.

Table 4. Hysteresis thresholds for classifying a node as fire-affected or smoke-affected.

Hazard Type	Entry Threshold	Exit Threshold
Fire	risk ≥ 0.45	risk < 0.30
Smoke	risk ≥ 0.25	risk < 0.15

The replanning trigger is activated immediately whenever a new fire or smoke node appears, or when an existing fire node clears; for gradually changing conditions, it is activated when the risk level changes by a value greater than or equal to 0.20, or the crowd density changes by a value greater than or equal to 0.15, for a given node. This stepwise, rather than continuous, response is the direct reason why, in six of the thirteen scenarios, the static fire and smoke node fields remain empty at the initial planning moment. The hazard itself enters the simulation solely through this sensor mechanism at a later stage (Section 4.4.3, Mechanism C).

The sensor confidence weights in Table 3 were set by engineering judgement of the authors, following a physically motivated ordering: temperature is the most direct indicator of active combustion and therefore receives the highest fire-weight; smoke obscuration is the most direct indicator of a tenability-impairing atmosphere and therefore receives the highest smoke-weight; CO_2 is treated as the least specific of the three, since its concentration can also rise from occupancy alone rather than combustion. No fire-safety standard prescribes a specific numerical scheme for fusing heterogeneous sensor types into a single hazard score: NFPA 72 [31] formally recognizes multi-criteria/multi-sensor detection as a detection strategy but explicitly delegates the mathematical evaluation criteria to the specific detection algorithm rather than specifying them normatively. The hysteresis thresholds in Table 4 are conceptually aligned with the tenability limits established in ISO 13571 [32], which defines compromised-tenability endpoints for heat, smoke obscuration, and toxic-gas exposure via the Fractional Effective Dose (FED) methodology; however, the present thresholds operate on the aggregated, normalized risk score rather than being derived through the formal

FED calculation, and closer alignment with this methodology is identified as a direction for future refinement (Section 6.6).

The numerical parameters of the hazard-assessment module (Tables 3 and 4, the normalization maxima and the severity coefficients of Equation 2) were set by engineering judgement, with ISO 13571 serving only as conceptual guidance and NFPA 72 prescribing no numerical values; unlike the fitted edge-cost weights, they determine which nodes are classified as hazardous and when, and thus act on the primary safety criterion, and their provenance and role in the ranking are classified in Section 3.9.

3.3. Building Connectivity Model

The topology of each building is represented as a weighted directed multigraph, in which nodes correspond to spatial elements (rooms, corridors, stairwells, exits) and edges correspond to physical transitions between them (doors, open spaces). The graph is described hierarchically (building → wing → floor → vertical connectivity), following the model from Ivanov [1], described in Section 4.1 in the context of the computational efficiency of the search.

Besides the dynamic exclusion of nodes depending on the current hazard, the graph also supports a permanent structural-exclusion mechanism, applied once at load time and independent of the simulation's progress. Nodes listed in the configuration as physically inaccessible are removed from the graph together with all edges connected to them. This mechanism is used for the regulatory exclusion of elevator nodes from the evacuation graph, since elevators cannot serve as an evacuation route during a fire unless specifically designed for that purpose. A separate, finer-grained mechanism allows the removal of a specific door or transition between two nodes without excluding the nodes themselves. This is suitable for doors that are regulatorily locked or permanently inaccessible, while the room itself remains part of the evacuation graph through other doors (a room is connected to adjacent rooms through multiple doors).

3.4. Route-Cost Model

Each node carries a composite risk score *risk_score* in the range $[0, 1]$, computed as a sum of indicator contributions, capped from above (Equation (2)):

$$risk_score(node) = \min(0.9 \cdot \mathbb{1}[fire] + 0.5 \cdot \mathbb{1}[smoke] + 0.3 \cdot \rho_{crowd}(node) + 1.0 \cdot \mathbb{1}[blocked], 1.0), \quad (2)$$

where $\mathbb{1}[\cdot]$ denotes the indicator function, equal to 1 if the condition in brackets holds for the node and 0 otherwise; $\mathbb{1}[\text{fire}]$ and $\mathbb{1}[\text{smoke}]$ indicate whether the node is currently classified as fire-affected or smoke-affected, respectively (Section 3.2); $\rho_{crowd}(node) \in [0, 1]$ is the node's normalized crowd density; and $\mathbb{1}[\text{blocked}]$ indicates whether the node is structurally impassable (e.g., a collapsed or obstructed link), independent of any sensor-detected hazard.

The coefficients reflect the relative severity of the hazards: a fire zone (0.9) is close to maximum risk, a smoke-affected zone (0.5) is a serious but potentially more briefly traversable hazard, and crowd density (0.3) adds a smooth, continuous penalty component rather than a binary prohibition. A combination of simultaneous fire and smoke reaches the maximum value $risk_score = 1.0$. This score is independent of the calibratable weight vector described below and plays a central role in the route-selection criterion (see Section 3.5).

As with the sensor confidence weights in Table 3, the coefficients in Equation (2) reflect the authors' engineering judgement of relative hazard severity, informed by, but not formally derived from, the tenability concepts underlying ISO 13571 [32].

For each edge of the graph, an eight-dimensional feature vector f is computed, summarized in Table 5, describing the conditions at the edge's target node.

Table 5. Eight-dimensional feature vector f computed for each edge of the graph.

Index	Feature	Semantics
f_0	bias	Constant base cost per hop (minimum number of steps)
f_1	fire_risk	Continuous fire risk score (or 0.9 for a binary indicator)
f_2	smoke_risk	Continuous smoke risk score (or 0.5 for a binary indicator)
f_3	crowd_density	Crowd density at the node, [0, 1]
f_4	blocked	1.0 if the connecting door is marked as blocked
f_5	exit_bonus	1.0 if the node is an exit
f_6	fire_proximity	1.0—normalized distance to the nearest fire node
f_7	distance	Geometric distance along the edge, normalized to 5 m

The cost of traversing an edge is computed as the dot product of the calibratable weight vector w and the feature vector f, increased by a contribution from the projected crowd density, and bound from below for the numerical stability of the search algorithm (Equation (3)):

$$edge_cost(node, link) = \max(w \cdot f(node, link) + \hat{c}(node) \cdot w_3,\ 0.1), \quad (3)$$

where f $(node,link) = (f_0, f_1, ..., f_7)$ is the eight-dimensional feature vector defined in Table 5, computed for the edge's target node (e.g., f_1 is that node's current fire_risk, f_6 its fire_proximity, and so on); $w = (w_0, w_1, ..., w_7)$ is the corresponding calibratable weight vector (see Table 6); and $\hat{c}(node)$ is the projected crowd density at the node, defined separately in Equation (4) below.

Table 6. Default weight vector for the edge-cost model.

Weight	Feature	Value	Direction of Effect
w_0	bias	1.0	Minimum cost per hop
w_1	fire_risk	10.0	Higher value → stronger avoidance of fire nodes
w_2	smoke_risk	5.0	Higher value → stronger avoidance of smoke-affected nodes
w_3	crowd	3.0	Higher value → stronger avoidance of congested nodes
w_4	blocked	15.0	Higher value → stronger avoidance of blocked doors
w_5	exit_bonus	−2.0	Negative value → attracts toward exit nodes
w_6	fire_proximity	5.0	Higher value → earlier, anticipatory avoidance of fire zones
w_7	distance	0.5	Higher value → preference for geometrically shorter paths

The default values of the weight vector w, used for the uncalibrated configurations in the ablation study (Section 4.3), are given in Table 6.

The projected crowd density $\hat{c}$(node), used in Equation (3), is obtained from a diffusion model that propagates observed density to neighboring nodes in the direction of the exits (Equation (4)):

$$\hat{c}(node, t+1) = c(node, t) \cdot \delta \cdot (1 - \varepsilon) + \bar{n}(node) \cdot \sigma \cdot (1 - \delta), \quad (4)$$

where $c(node,t) \in [0, 1]$ is the projected crowd density at the node at time step t; $\bar{n}(node)$ is the density averaged over the node's graph neighbors, weighted by their distance to the nearest exit, so nearer-to-exit neighbors count more—making the diffusion directional rather than uniform. At each step, the node's own density decays by retention $\delta = 0.85$ (with a further $\varepsilon = 0.15$ evaporating), while a $\sigma = 0.30$ share of the neighbors' density flows in. The result is then blended with the latest direct sensor reading via exponential smoothing ($\alpha = 0.4$), so a real measurement always overrides the diffused estimate. In this way, nodes closer to the exit have greater influence, modeling directed rather than isotropic diffusion.

Direct sensor observations are given precedence over the projection through exponential smoothing with coefficient $\alpha = 0.4$: $c(node)0.4observed(node) + 0.6c(node)$.

3.5. Route Search and Selection

During the search for K alternative routes, nodes are traversed via a priority queue (min-heap) whose key reflects not only the accumulated path cost but also current risk conditions and the accumulated long-term experience of each node's quality (Equation (5)):

$$priority(node,cost) = cost - 0.5 \cdot P(node) + 8.0 \cdot risk_score(node) + 4.0 \cdot \rho_{crowd}(node) + \Delta_{up}(node), \quad (5)$$

where the first term $P(node)$ is the node's accumulated long-term value, the second $risk_score(node)$ is the composite risk score from Equation (2) (with a coefficient of 8.0, which imposes a strong penalty contribution for hazardous nodes, independent of the calibratable weight w_1), the third 4.0 $\rho crowd(node)$ adds a penalty for overcrowded nodes, and the fourth $\Delta up(node)$ imposes a penalty of 12.0 for every floor above the agent's starting floor, which prevents routes that climb back upward without good reason.

The number of alternative routes searched at each planning cycle is bound by parameter K. For the scenarios evaluated in Section 4, the value applied is $K = 5$ for every planning cycle, regardless of the value specified in each individual scenario's configuration. A direct comparison of the results for $K = 5$ versus $K = 200$ shows that, for twelve of the thirteen scenarios, the number of evacuated agents remains completely identical, and in the single scenario with a measurable difference, the deviation is one agent. This comparison is limited to the ranking criterion and conditions of Section 6.5 and is not a general result.

After the K admissible routes are found, they are sorted using a two-criterion lexicographic key function (Equation (6)):

$$sort_key(n) = (risk_level(n), total_cost(n)), \quad (6a)$$

$$risk_level(n) = \frac{1}{n}\sum_{i=1}^{n} risk_score(n_i), \quad (6b)$$

where the term $total_cost(n)$ is the sum *edge_cost* over all n edges of the path. The semantics of this sorting are central to the algorithm's architecture: the safer route (lower value *risk_level*) is always preferred over the cheaper one, regardless of the difference in the secondary term *total_cost*. The calibratable weight vector w participates only in the secondary criterion *total_cost* and therefore decides only between routes with an identical *risk_level* primary value. The route with index 0 after sorting is the recommended route for the corresponding agent. This lexicographic structure is called TWO_TIER.

It is important to state precisely what this order guarantees. Nodes classified as fire- or smoke-affected (Section 3.2) are removed from the graph before the search (Section 5.2), so no candidate route can contain them, irrespective of any weight. Among the remaining candidates, the primary key *risk_level* depends only on the fire, smoke and blocked indicators and on crowd density (Equation (2)); proximity to a fire node does not enter it. Two routes with equal *risk_level* may therefore differ in their distance to the nearest classified fire node, and the choice between them is made by *total_cost*, in which proximity enters only through the term $w_6 f_6$ (Table 6). With the default positive w_6 this term penalizes proximity, but this is a property of the weight values and not an architectural guarantee: it can be outweighed by the bias and distance terms of a shorter route, and the offline calibration (Section 3.6) imposes no sign constraint on w_6. Likewise, because crowd density contributes to *risk_level*, a route closer to an unclassified fire front may be ranked first by the primary key if it is less congested.

3.6. Weight Calibration and Online Adaptation

The weight vector w from Table 6 can be loaded from a pre-calibrated model instead of using the default values. Calibration is performed offline, via gradient descent, minimizing the mean-squared error between the cost predicted by Equation (3) and an "ideal" cost derived from simulated evacuation trajectories in the test scenarios. This linear model with eight defined features contains no graph neural network, no message-passing mechanism between nodes, and no learned node embeddings. Calibration is limited to the eight scalar coefficients of the linear combination in Equation (3).

Alongside the offline-calibrated weight vector, the system maintains an online table of long-term node values $P(node)$, used in Equation (5). Unlike formal reinforcement learning, this table uses no reward function, discount factor, or Bellman equation—it is updated through direct, threshold-based addition of fixed correction values whenever certain events occur, summarized in Table 7.

Table 7. Threshold-based corrections to the online table of node values $P(node)$.

Event	Correction to $P(Node)$
The node is in a fire zone	−15.0
The node is in a smoke-affected zone	−5.0
The node is a free exit	+10.0
The node is overcrowded	−3.0

These corrections act as a long-term memory of node quality: after a series of successful evacuations, nodes that were part of reliable routes accumulate a positive value and are preferred earlier in subsequent searches, without this constituting a formal reinforcement-learning process in the sense of a Markov decision process.

As shown in Section 4.4, these values are not safety-critical, since route safety is governed independently by the hard-coded *risk_level* term.

3.7. Coordination Among Agents

When planning for multiple agents, the processing order is not arbitrary: more endangered agents receive a computed route first, before the virtual load from already-planned agents can influence their choice (Equation (7)):

$$priority(agent) = \begin{cases} 3.0 & \text{if the agent is in a fire or an area with unbearable smoke} \\ 2.0 & \text{if the agent is in a tolerable smoke} - \text{filled area} \\ 1.0 + \rho_{crowd}(agent) & \text{otherwise} \end{cases}, \quad (7)$$

Under the strict smoke-tolerance threshold configuration used in the evaluated scenarios (Section 5.2), every classified smoke node is treated as intolerable, so that in practice an agent in a smoke node receives the same highest priority (3.0) as an agent in a fire node. A value of 2.0 would apply only under a looser tolerance threshold, permitting low smoke levels, different from the configuration used in the present study.

To prevent multiple agents from concentrating on the same route, after each agent is planned, the nodes along its route receive a temporary, virtual increase in crowd density (Equation (8)):

$$working_density(node) \leftarrow \min(working_density(node) + \lambda, 1.0). \quad (8)$$

The value of the coefficient λ applied in all simulation runs in Section 4 (the PROPOSED configuration) is $\lambda = 0.35$. Other values of λ were not evaluated in the main experiments, so all coordination results reported below are valid only for this setting (see Section 6.3).

The next agent in sequence "sees" the already-planned nodes as more congested and naturally prefers an alternative route, without any need for centralized coordination or communication among agents. This mechanism is quantitatively confirmed in Section 4.5 through the ablation configuration NO_AGENTS_COORD ($\lambda = 0$).

When an agent is left without an admissible route (*has_solution* = *false*) for two or more consecutive planning cycles, a rescue-intervention notification is automatically generated, containing the agent's current position, a recommended access point for rescue teams, and the duration of the blockage. This sequence is directly verified in the simulation log for the boundary case.

3.8. Crowd-Movement Model

The route computed in Section 3.5 is not executed unconditionally. The physical movement of each agent from node to node is governed by a separate, independent model, implemented as a discrete cellular automaton with a fixed, at this stage, five-second time step. This model acts as an execution layer between the planned route and the agent's actual position, and includes its own, independent hazard-avoidance logic, distinct from the hard exclusion of nodes at the planning level. This is a defence-in-depth architecture, in which physical entry into a fire or intolerable-smoke node is prevented at two independent levels.

The agent's movement speed depends on the local crowd density at the node the agent currently occupies (Equation (9)):

$$speed(agent) = \begin{cases} mobility(agent) & \text{if} congestion < 0.3 \\ \frac{mobility(agent)}{1+(congestion-0.3)\cdot 3.0} & \text{otherwise} \end{cases}. \quad (9)$$

Speed accumulates as credit at each step. The agent changes node only once the accumulated credit reaches 1.0, after which the credit is reduced by 1.0. This mechanism is designed to allow agents with different individual mobility to advance at different effective rates within an equal discrete time step; in the present evaluation, however, all agents were assigned the default mobility value (1.0), so this capability was not exercised in the reported scenarios.

The decision as to which neighboring node is selected as the next position follows a three-level priority hierarchy. First, if the agent is already located in a fire or intolerable-smoke node, escape logic is activated: the agent follows the route recommended by the planner if it points to a neighboring node that is not fire-affected or intolerably smoke-affected. Otherwise, a neighboring node is chosen via a scoring function that gives preference to fully safe nodes over tolerably smoke-affected nodes, and to fire nodes only as a last resort, with an additional penalty for recently visited nodes to prevent persistent oscillation between two neighboring nodes. Second, if the agent is not in danger and a planned route exists, it is followed to the next node along the path, unless that node is fire-affected or intolerably smoke-affected, in which case processing moves to the third level. Third, in the absence of an applicable planned move, direct, local (greedy) selection among the neighboring nodes is applied, based on an attraction scoring function (Equation (10)):

$$attraction(v) = w_{exit} \cdot exit_score(v) - w_{danger} \cdot danger_score(v) - w_{crowd} \cdot crowd_density(v), \quad (10)$$

with default weights *wexit* = 2.0, *wdanger* = 3.0, *wcrowd* = 1.0;the first term $exit_score(v)$ is a monotonically decreasing function of the number of hops to the nearest exit (obtained via multi-source breadth-first traversal from all exit nodes), and the second term $danger_score(v)$ is 1.0 for a fire node, 0.5 for a smoke node, and 0.0 for a safe node. At all three levels of the hierarchy, a choice leading to a fire or intolerable-smoke node from a currently safe position is explicitly prohibited and is confirmed through a separate validation step after all moves for the current step have been determined, before they are applied.

If all an agent's neighboring nodes are simultaneously on fire, the agent remains at its current position rather than being forced to choose among equally inadmissible options. This is a state corresponding to a section of the graph topologically cut off by the hazard, and it is the threshold condition for the rescue-intervention notification. It is precisely the discrete, stepwise nature of this movement model that explains why an agent already located in a given node at the moment of its sensor classification as fire- or smoke-affected may register brief physical contact with the hazard before the escape mechanism has the opportunity to trigger at the next step.

The parameters of the movement layer are the five-second time step; the speed rule of Equation (9) (congestion threshold 0.3, slope 3.0); the attraction weights of Equation (10) ($w_{exit} = 2.0$, $w_{danger} = 3.0$, $w_{crowd} = 1.0$); the node capacities of Table 2; and the homogeneous mobility 1.0. These values were set by engineering judgement and were not calibrated against empirical pedestrian data, and capacity is enforced per node and not per door width.

3.9. Parameter Provenance and Role in the Primary Ranking

The numerical parameters of the model fall into four categories. (1) Offline-fitted: the eight edge-cost weights of Table 6, which enter only the secondary criterion *total_cost* (Equation (6a)). (2) Engineering judgement: the sensor weights (Table 3), hysteresis thresholds (Table 4), normalization maxima and the severity coefficients of Equation (2). (3) Regulatory or standard-derived: the zero smoke tolerance and the exclusion of elevators (Regulation Iz-1971) and the scenario inputs of Section 4.1 (fire growth, occupant loads, pre-movement times). (4) Uncalibrated design constants: the coefficients of Equation (5) and Table 7, the capacities of Table 2, the diffusion constants of Equation (4), the replanning triggers, λ, K and the time step; these were neither fitted nor derived from a standard, and only K, λ (exploratorily) and the sensing interval (Section Sensor-Interval Sensitivity and Preconditions of the Hazard-Avoidance Property) were examined. The safety-governing primary key *risk_level* depends on categories 2 and 3 and, through the crowd-density term of Equation (2), on the diffusion constants, capacities and λ of category 4, but not on the fitted weights of category 1. Categories 2 and 4 therefore carry the model's safety-relevant assumptions and should be validated before operational use. The fitted edge-cost weights of category 1 cannot re-admit a node that has already been classified as hazardous. Changing the parameters of category 2 alters the moment and extent of node exclusion and therefore the exposure of agents, in a way that is not necessarily monotonic (Section 6.6).

4. Experiments

The results of Sections 4 and 5 are held under the following hypotheses. (1) Evaluation is by simulation only; absolute evacuation times are not real-world predictions. (2) Every agent receives its route instantaneously and follows it with full compliance; the only departures from the plan are the movement rules of Section 3.8. (3) Sensing is ideal: the sensor series are deterministic, complete, noise-free and sampled every five seconds; sensor failure, false alarms and missed alarms are not modelled. (4) The hazard is prescribed by the scenario sensor series and is not affected by the evacuees. (5) Agents are homogeneous (mobility 1.0, speed depending only on local crowd density, no smoke-dependent slowing). (6) The sensor weights, hysteresis thresholds, severity coefficients and coordination coefficient $\lambda = 0.35$ are fixed and were set by engineering judgement. (7) The regulatory analysis refers to a single jurisdiction. (8) The comparisons are with ablated variants of the proposed system only. Limitations arising from these assumptions are discussed in Section 6.6.

4.1. Experimental Setup

The algorithm was evaluated in a simulation environment combining (i) an agent-based crowd-movement model, implemented as a cellular automaton, and (ii) a risk-aware K-shortest-paths planner, which computes K shortest paths over a building connectivity graph weighted by a linear edge-cost model with eight defined parameters (presence of fire, presence of smoke, crowd density, blocked connections, proximity to exit, proximity to fire, geometric distance, constant offset value). The weight vector is calibrated offline via gradient descent on a mean-squared-error-type loss function relative to an "ideal" cost derived from simulated trajectories and is further adapted online through threshold-based updating of a table of node values as new sensor data arrives.

Building topology is described through a hierarchical directed-multigraph connectivity model [1], in which nodes represent spatial elements (rooms, corridors, stairwells) and edges are the doors and openings between them. Unlike the flat graph-based models used in most indoor-navigation systems, the hierarchical representation (building → wing → floor → vertical connectivity) allows a subgraph to be extracted for only the relevant part of the building, without loading the entire topology into memory. This advantage was demonstrated experimentally in the original paper (up to 1433 times faster than Yen's algorithm on complete graphs, and 4.7–27 times faster in real building scenarios, relative to the same baseline algorithm). Provision is made for introducing virtual nodes. For example, a virtual target node (outdoors) allows simultaneous search toward multiple admissible exits in a single call to the algorithm.

Hazard (fire and smoke) is modeled dynamically through time series of sensor data, which are fed into the simulation from a file in a stepwise manner and lead to recomputation of the plans for affected agents whenever a significant change in hazard occurs, rather than only once at the start of the simulation.

The fire and smoke dynamics in Scenarios 10–13 are constructed to reflect established fire-safety engineering practice, rather than arbitrarily chosen curves: the growth of heat-release rate follows the t^2 law [33,34], with a growth class (ultra-fast/fast/medium) determined by the space category; the temperature contribution at each node is computed via the Alpert [35] ceiling-jet correlation; occupant loading per area follows BS 9999 [36] and the occupancy load tables of NFPA 101 [37]; pre-movement times follow PD 7974-6 [38] only where they are modelled.

The evaluation covers thirteen scenarios, presented in Table 8, spanning four different buildings and differing in the number of agents, topology, and fire-spread dynamics.

The simulation advances in five-second steps and runs for at most three times the length of the sensor series. The sensor series are deterministic, with no noise added to the readings, and are sampled every five seconds. The Monte Carlo replicates differ only in two random elements: a random integer start-time jitter of 0–10 s for every agent, and a uniform perturbation of ± 0.08 of the initial crowd densities (clipped to $[0, 1]$) at the first planning step.

Scenarios 1–7 use manually defined hazard series (3 to 32 agents); start delays of up to 55, 40 and 30 s are used in Scenarios 1, 3 and 4, and all other agents start at $t = 0$. In Scenario 8, 64 agents (24, 24 and 16 on floors 4, 3 and 2) all start at $t = 0$, and a designed fire in the west stairwell is classified as fire at 10, 30 and 50 s in successive stairwell segments, with smoke reaching the main corridors of floors 2 and 3 at 45 and 60 s. In Scenario 9, 80 agents are distributed over five floors (12, 16, 22, 22 and 8) and all start at $t = 0$; a slowly growing fire in the second-floor corridor is classified as fire at about 50 s.

The ZDB scenarios use t^2 fire growth with Alpert's ceiling-jet correlation. In Scenario 10 (48 agents, one per room) three vehicles ignite at 0, 300 and 600 s with ultra-fast growth ($\alpha = 0.1874$ kW/s^2, Qmax = 6000 kW, ceiling height 2.7 m). Scenarios 11–13 (120, 120 and

80 agents) use area-based occupancy and growth classes ultra-fast (75 s), fast (150 s) and medium (300 s) according to space use; Scenarios 12 and 13 add car-to-car spread with a 25 s ignition delay, and Scenario 13 additionally contains a stairwell link that is blocked from the start.

Table 8. Overview of the thirteen evacuation scenarios used for the experimental evaluation.

No.	Building	Agents	Nature of the Hazard
1	Rectorate, 3rd floor	5	Cascading evacuation, fire in the stairwell
2	Rectorate, floors 2–4	8	Simultaneous mass evacuation
3	Library, floors 1–3	6	Cascading evacuation with smoke
4	Rectorate + Library	6	Mixed evacuation from two buildings
5	Rectorate	3	Blocked agent, fire-department notification
6	Rectorate + Library	32	Mass planned evacuation
7	Rectorate + Library	32	Mass evacuation under stress
8	Rectorate	64	Expanding fire (floors 2/3/4)
9	Institute	80	Congestion at a single connectivity core
10	ZDB	48	Fire spread from several vehicles
11	ZDB	120	CFD-type fire (area-based modeling)
12	ZDB	120	CFD-type fire, car-to-car
13	ZDB	80	Cascading fire with stairwell blockage

Pre-movement times are modelled only in some scenarios. In Scenario 10 they are log-normal (median 30 s, $\sigma = 0.4$ for garage occupants; median 120 s, $\sigma = 0.5$ for apartments; truncated at 480 s). In Scenarios 12 and 13 parking occupants (34 of 120 and 40 of 80 agents) start after fixed delays of 20–40 s. In Scenarios 8, 9 and 11 all agents start at $t = 0$.

Detailed information for each scenario (objective, building(s), number of agents, fire type and dynamics, smoke type and dynamics) is available in Supplementary Material/docs/SF2_Software_description.pdf.

The Rectorate and University Library buildings are real buildings of the Technical University of Gabrovo (Bulgaria), modeled specifically for the purposes of the present study. The Institute corresponds to a publicly available IFC test model [39], and ZDB corresponds to a real test model for evacuation-route analysis, published by the research team at TU Wien [40] under a Creative Commons Attribution 4.0 license.

A visual overview of the four studied buildings, illustrating the difference in their scale and typology, is shown in Figure 2. For the Institute and for ZDB individually, the corresponding building IFC file serves as the sole source of information—both the 3D geometry and the connectivity graph used by the routing algorithm were extracted from it. For the Rectorate and the Library, the 3D visualization and the connectivity graph are two separate representations, independent of one another, of the same real buildings: the 3D model was built in SketchUp solely for illustration purposes, while the connectivity graph was defined manually.

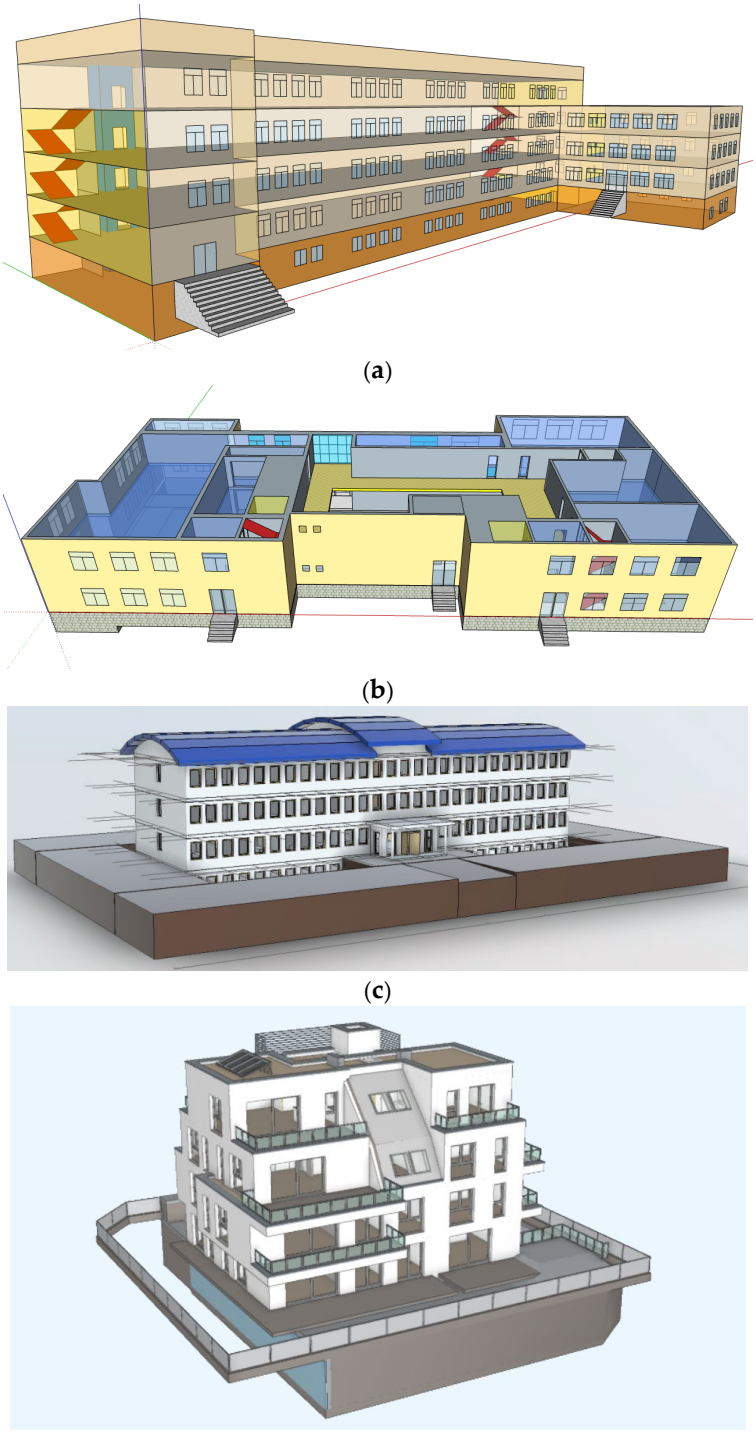


Figure 2. 3D views of the four studied buildings: (**a**) Rectorate, (**b**) University Library, (**c**) Institute, (**d**) ZDB.

4.2. Evaluation Metrics

Six metrics were used for the quantitative evaluation of the algorithm, computed from the complete agent-movement log rather than from the planned routes alone, to reflect the executed behavior.

Evacuated (number of evacuated agents). The number of agents that reach the "outdoors" node by the end of the simulation, relative to the total number of agents in the

scenario. This information is complemented by Blocked agents—agents for whom the planner found no admissible route and for whom the emergency services are notified.

Mean and maximum evacuation time. The mean and maximum time (in seconds) required successfully evacuated agents to reach an exit. The maximum value is indicative of the worst-case scenario. This is a critical metric for compliance with regulatory evacuation-time requirements.

Number of exits used. The number of distinct exits through which agents left the building, relative to the total number of available exits in the scenario. A low value indicates overloading of a small subset of exits, even given a good mean evacuation time.

Gini coefficient of exit load. Computed according to Equation (11), where the values v_i are the numbers of evacuated agents per exit, sorted in ascending order, and n is the number of exits actually used—with one correction: if all evacuated agents used a single exit ($n = 1$) while the building has more than one registered exit ($N_{door} > 1$), n is instead set to N_{door} and the count vector is zero-padded with the $N_{door} - 1$ unused exits before evaluating the sum below. This distinguishes a genuine single-exit building, for which $G = 0$ is the only possible outcome, from a case where several exits exist but agents happened to concentrate on only one of them.

$$G = \frac{2\sum_{i=1}^{n} i \cdot v_i}{n \cdot \sum_{i=1}^{n} v_i} - \frac{n+1}{n}. \tag{11}$$

The value of $G = 0$ corresponds to perfectly uniform loading of all exits; the maximum value ($G \to 1$) corresponds to complete concentration of evacuation on a single exit. A building with only one registered exit trivially yields $G = 0$; by contrast, all agents using only one of several available exits is scored close to $G = 1$. This metric is key for assessing the contribution of inter-agent coordination (Section 4.5), since it directly reflects load balance rather than evacuation speed.

Safety-compliance metrics (agents_fire/agents_smoke). The number of evacuated agents whose executed movement path (rather than the planned route) passed through at least one node classified as fire-affected, or respectively smoke-affected, at the moment the agent occupied it. This is computed by matching each agent's position log against the sensor time series for fire and smoke. This metric is independent of the speed and balance metrics above and is used specifically to verify regulatory compliance. These two metrics cover evacuated agents only; non-evacuated agents are reported separately through their rescue notifications and tenability margins.

We tested differences with the two-sided Mann-Whitney U test. We measured their size with the signed Cliff's delta. For time metrics, a positive value means the proposed system took longer. We made 260 comparisons in total, and we treated them as a single family to account for multiple testing. For this we applied the Holm-Bonferroni and Benjamini-Hochberg adjustments. Unadjusted p-values and all results for the small scenarios should be regarded as exploratory.

4.3. Ablation Configurations

To isolate the contribution of each individual component of the proposed algorithm, an ablation study was conducted with five configurations, distinguished by four independent switches: (1) whether the weight vector is offline-calibrated or set manually to its default values; (2) whether online adaptation of the node-value table is active; (3) the degree of coordination among agents in distributing load across exits; and (4) whether the plan is dynamically recomputed upon a significant change in hazard, or computed once at the start. The five configurations are summarized in Table 9.

Table 9. Ablation configurations used in comparative study.

Configuration	Calibrated Weights	Online Adaptation	Coordination	Dynamic Replan
PROPOSED	yes	yes	0.35	yes
STATIC_WEIGHTS	no	no	0.35	yes
NO_AGENTS_COORD	yes	yes	0.0	yes
KSP_BASELINE	no	no	0.0	yes
STATIC_PLAN	no	no	0.0	no

PROPOSED is the complete version of the proposed algorithm. STATIC_WEIGHTS isolates the effect of offline calibration and online adaptation on its own, keeping coordination and dynamic replanning unchanged. NO_AGENTS_COORD isolates the effect of coordination on its own, keeping calibration and adaptation active. Here the coordination factor is set to zero, so that each agent plans its route independently of the load on the exits chosen by the other agents. KSP_BASELINE combines both restrictions but retains dynamic replanning. It simulates a standard K-shortest-paths planner with a fixed heuristic, extended solely with a response to hazard changing over time. STATIC_PLAN isolates the contribution of dynamic recomputation on its own. In this case, the plan is computed once at the start of the simulation and is not subsequently updated, regardless of the later development of the fire or smoke.

Each of the five configurations was run with twenty Monte Carlo replicates for each of the thirteen scenarios, using identical, paired values of random jitter and density-prediction noise across configurations for a given replicate index. This is a methodological choice that reduces variance in the comparison and allows the observed differences to be directly attributed to the factor under study.

4.4. Sensitivity Analysis

4.4.1. Methodology

To verify the extent to which the proposed linear edge-cost model governs route selection, a One-At-a-Time (OAT) local sensitivity analysis was conducted on the eight calibratable weights, following the methodology recommended in ASTM E1355 [41] for assessing the predictive capacity of deterministic fire models. For each weight, its value is scaled relative to the default value by six factors (0.25, 0.5, 1, 1.5, 2, and 3) while the remaining seven weights are held fixed at their default values. The base case is shared across the sweeps of the individual weights, so the total number of runs per scenario is 41. The analysis was applied across all thirteen scenarios to verify whether the conclusions are specific to a particular topology or generalize (Supplementary Material/results/sensitivity).

4.4.2. Results

The results, summarized in Table 10, show that a substantial share of the eight weights have no measurable effect on the output metrics across the entire range studied.

Table 10. Number of weights with no measurable effect ("flat" weights) out of eight total, per scenario.

Scenario	Flat Weights	Which Weights
1–5, 7	8/8	bias, fire, smoke, crowd, blocked, exit, fire_proximity, distance
6, 8, 9, 10	5/8	fire, smoke, blocked, fire_proximity, distance (for Scenario 9: +exit)
11, 12, 13	6/8	fire, smoke, blocked, exit, fire_proximity, distance

Flatness is not confined to a particular scenario. The weights fire, smoke, and blocked are systematically flat across all scenarios, regardless of the fire's characteristics. For example, Scenario 8, in which the fire spreads across three floors, shows an identical flatness pattern to Scenario 10, in which the fire is localized.

4.4.3. Causal Analysis

To determine whether the lack of sensitivity is an implementation defect or architecturally expected behavior, the planner was directly instrumented to trace what happens during each route-planning step. Three distinct, empirically confirmed mechanisms were identified.

Mechanism A: absence of an alternative route. In the scenarios with limited connectivity (1–5, 7), only a single physically feasible route to the exit exists for most agents. A direct test with the fire weight ranging from 2.5 to 30.0 shows an identical total route cost and an identical selected route, since the sole feasible path never passes through a node with a nonzero value for the corresponding feature.

Mechanism B: dominance of a fixed risk-gate over the calibratable cost. The final route selection sorts of candidate routes by a dual key (*risk_level*, *total_cost*), where the primary term *risk_level* is a separate, hard-coded formula, entirely independent of the eight calibratable weights. The latter enter only into the secondary term *total_cost*, which serves solely as a tie-break criterion. Direct instrumentation of all planning points in Scenario 10 (869 agent-replan instances) shows that changing the fire weight from 10.0 to 30.0 changes the total cost in 108 of 869 cases (12%) but never changes the selected route (0 of 869 cases).

Mechanism C: dynamic rather than static fire modeling. The static configuration fields for fire and smoke nodes are empty in all six of the larger scenarios (8–13); fire and smoke enter only dynamically, through the sensor time series, later in the simulation's timeline. The initial plan, computed at the starting instant $t = 0$, is computed under an entirely clean hazard state. The fire and smoke parameters are structurally zero at this planning point, regardless of the value of the corresponding weights.

An illustration of the mixed sensitivity pattern is presented in Figure 3, showing the results for Scenario 10: the panels for fire, smoke, blocked, fire_proximity, and distance are completely flat, while the panels for crowd and bias show a measurable dependence.

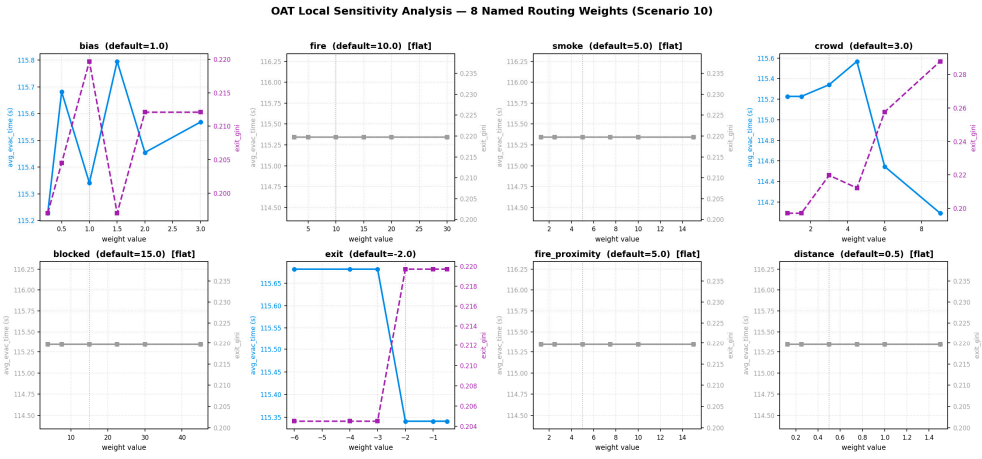


Figure 3. OAT local sensitivity analysis for the eight calibratable weights, Scenario 10. In each panel, one weight is varied while the other seven are held at their default values; the default value is given in the panel title and marked by the vertical dotted line. The solid blue line with circles (left axis) shows the average evacuation time (avg_evac_time, s), and the dashed purple line with squares (right axis) shows the Gini coefficient of exit load (exit_gini). Grey lines, with panels marked [flat], indicate that neither metric changed over the entire range studied, i.e., no measurable effect of the weight.

4.4.4. Interaction with the Compensating Regulatory-Compliance Correction

The hard exclusion of fire and smoke nodes from the planning graph (introduced for regulatory compliance, Section 5) has a side effect on the sensitivity-analysis results. The reason is a logical consequence of Mechanism A: removing previously generated inadmissible routes reduces the total number of alternative candidate routes, and hence the tie-break opportunities through which these two weights had shown an effect. This is a quantitative picture of the architecture, paid as a necessary price for compliance with regulatory requirements for fire evacuation from buildings.

4.4.5. Scope of the Tie-Break and Proximity to Fire

The lexicographic ordering guarantees the exclusion of classified hazard nodes but, by construction, does not guarantee that the route chosen among equally risky candidates is the one farthest from the fire. Direct instrumentation across all thirteen scenarios identified 965 tie-break events, in none of which a candidate with worse fire proximity was selected (Section 6.6). This observation is empirical: with default weights, $w_6 > 0$ biases the tie-break away from fire, and tied candidates tended to have near-identical proximity. We therefore do not claim a formal guarantee. Two constraints are recommended for any deployment: (i) the calibration of w_6 should be constrained to non-negative values (or w_6 fixed at its default), so that the proximity term can never reward proximity to fire; this does not by itself prevent a shorter route from being chosen over a farther one; (ii) where proximity to unclassified fire fronts is safety-relevant, it should be raised into the primary criterion or enforced by a hard buffer around active fire nodes. The latter was tested in this study (1-hop buffer) and altered outcomes in 4 of 13 scenarios, with up to a 67% increase in maximum evacuation time in one case, so it is not cost-free and is left to a dedicated development effort involving a certifying fire-safety engineer.

4.4.6. Conclusions

The observed low sensitivity is a consequence of the design decision to use a safety-oriented route-selection criterion. In this case, the calibratable weights act solely as a secondary mechanism with limited scope of influence. This is because these eight coefficients enter only the secondary tie-break term of the lexicographic sort key (Equation (6a)), while exclusion of classified hazard nodes is governed by the hard exclusion and the primary key *risk_level*. Varying each weight between 0.25× and 3× its default value produces no change in the selected route for most tested cases, confirming that their exact numerical values, whether set to the defaults in Table 6 or offline-calibrated, are not critical for the avoidance of classified fire and smoke nodes, although they may influence the choice among equally risky routes. This finding directly motivates the content of Section 4.5, in which the comparative Monte Carlo study localizes the real source of result improvement outside weight calibration.

The fixed coefficients of Equation (2), which enter the primary key *risk_level*, were also examined for sensitivity: each of the fire (0.9), smoke (0.5) and crowd (0.3) coefficients was varied by ±20% one at a time (Scenarios 10–12, five replicates). The numbers of evacuated agents, the maximum evacuation times, the contacts with fire or smoke and the rescue notifications were unchanged, and the exit-load Gini coefficient changed by at most 0.006. This is expected from the architecture: classified fire and smoke nodes are hard excluded from the candidate routes, so the fire and smoke terms can act only through the starting node common to all candidates, and the crowd term scales the mean density of a route.

The architectural question of whether a different choice of primary sorting criterion could reveal broader influence of the calibratable weights, without compromising the

hazard-avoidance property, remains an open architectural question, discussed further in Section 6.7.

4.5. Comparative (Monte Carlo) Study

4.5.1. Overview of Results

The aggregated result at the level of the entire sample of thirteen scenarios across five configurations and twenty replicates, a total of 1300 simulation runs and 260 pairwise comparisons (Supplementary Material/results/mc), is presented in Table 11.

Table 11. Number of scenarios (out of 13 total) with $p < 0.05$ (Mann–Whitney U) when comparing PROPOSED against each baseline configuration, by metric.

Metric	vs. STATIC_WEIGHTS	vs. NO_AGENTS_CO-ORD	vs. KSP_BASELINE	vs. STATIC_PLAN
Evacuated (count)	0/13	2/13	2/13	2/13
Mean evac. time	1/13	9/13	9/13	11/13
Max. evac. time	0/13	7/13	6/13	9/13
Number of exits	0/13	8/13	9/13	11/13
Gini coefficient	3/13	8/13	8/13	9/13

Without correction, 114 of the 260 comparisons are significant. After the Benjamini-Hochberg adjustment, 110 remain significant, and after the Holm adjustment, 79 remain. Two sets of results hold up after correction: the number of exits used and the Gini coefficient, both compared with the configurations without coordination. For the Gini coefficient, the direction was not always in favor of coordination. For example, it was higher in Scenarios 4 and 8. The results on mean evacuation time and the number of evacuated agents does not hold up after correction.

The PROPOSED configuration versus STATIC_WEIGHTS is statistically indistinguishable in practice almost everywhere. This is a direct confirmation of the finding from Section 4.4 (Mechanism B). PROPOSED versus configurations without coordination or without dynamic replanning shows a systematically significant effect. This aggregated result, however, conceals substantial variation among scenarios, which cannot be understood without a breakdown by group, presented in Figure 4 and the following subsections.

For a meaningful interpretation, the scenarios are divided into three groups based on two objective criteria: (1) the source of graph connectivity—manually described (Scenarios 1–8, the Rectorate and Library buildings) versus automatically extracted from IFC (Scenarios 9–13, the Institute and ZDB buildings); and (2) the number of actually accessible exits, which determines whether the load-balance metrics are structurally informative.

4.5.2. Group A: Small Buildings, Manual Connectivity, Limited Exit Choice (Scenarios 1–5)

The results for Group A are summarized in Table 12.

The number of exits used did not decrease in any scenario of this group, while the Gini coefficient was lower with coordination in some scenarios (2 and 5) and higher in another (4). A likely reason is that the sample is structurally limited: with 3–8 agents and 20 replicates per configuration, the distribution of "which agent through which exit" has too few possible outcome states for the Mann–Whitney test to reject the null hypothesis consistently. The effect sizes in Table 12 are comparable in magnitude to those in the larger groups, but because the p-values are unstable, the results for this group are regarded as exploratory.

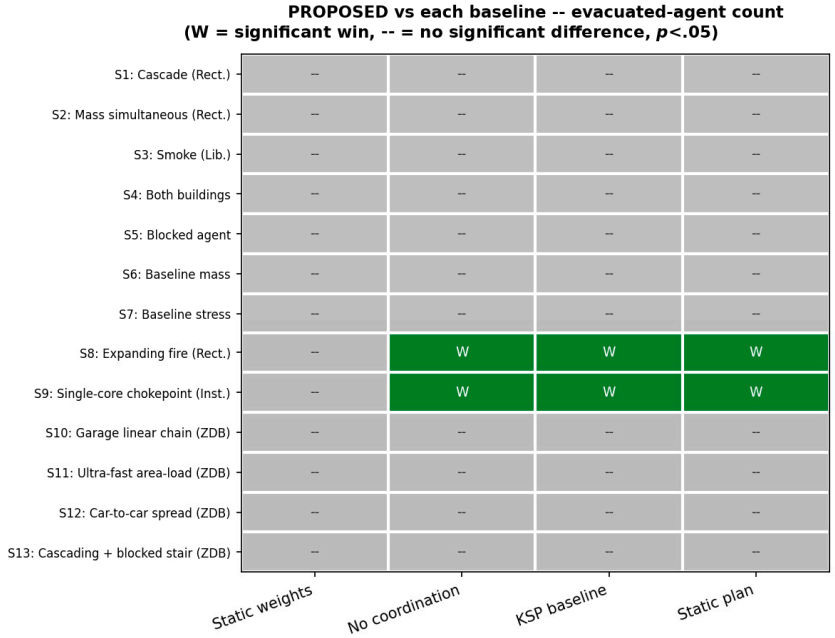


Figure 4. PROPOSED versus each baseline configuration by number of evacuated agents, for all thirteen scenarios (W = statistically significant win, – = no significant difference, $p < 0.05$).

Table 12. Significance (number of scenarios out of 5, $p < 0.05$) and mean effect size (|Cliff's delta|) for Group A.

Metric	vs. STATIC_WEIGHTS	vs. NO_AGENTS_COORD	vs. KSP_BASELINE	vs. STATIC_PLAN
Mean evac. time	0/5 ($\delta = 0.01$)	3/5 ($\delta = 0.33$)	3/5 ($\delta = 0.33$)	4/5 ($\delta = 0.43$)
Number of exits	0/5 ($\delta = 0.00$)	3/5 ($\delta = 0.64$)	3/5 ($\delta = 0.64$)	4/5 ($\delta = 0.80$)
Gini coefficient	0/5 ($\delta = 0.02$)	3/5 ($\delta = 0.65$)	3/5 ($\delta = 0.65$)	4/5 ($\delta = 0.80$)

4.5.3. Group B: Larger Buildings, Manual Connectivity, Rich Exit Choice (Scenarios 6–8)

The results for Group B, summarized in Table 13, show the strongest and statistically cleanest effect of the entire dataset.

Table 13. Significance and mean effect size for Group B.

Metric	vs. STATIC_WEIGHTS	vs. NO_AGENTS_COORD	vs. KSP_BASELINE	vs. STATIC_PLAN
Mean evac. time	1/3 ($\delta = 0.29$)	3/3 ($\delta = 1.00$)	3/3 ($\delta = 0.98$)	3/3 ($\delta = 0.90$)
Max. evac. time	0/3 ($\delta = 0.02$)	2/3 ($\delta = 0.71$)	2/3 ($\delta = 0.68$)	3/3 ($\delta = 0.91$)
Number of exits	0/3 ($\delta = 0.00$)	2/3 ($\delta = 0.67$)	2/3 ($\delta = 0.67$)	3/3 ($\delta = 1.00$)
Gini coefficient	2/3 ($\delta = 0.55$)	3/3 ($\delta = 1.00$)	3/3 ($\delta = 1.00$)	3/3 ($\delta = 0.80$)

Cliff's delta reaches a value of 1.00 (complete separation of the distributions) for the Gini coefficient and the number of exits relative to the three uncoordinated configurations. The reason is direct: with 32–64 agents and 7–8 accessible exits in a richly connected topology, the coordination mechanism has enough agents and alternative exits available to display its full potential for load balancing. An illustration for the representative Scenario 6 is shown in Figure 5.

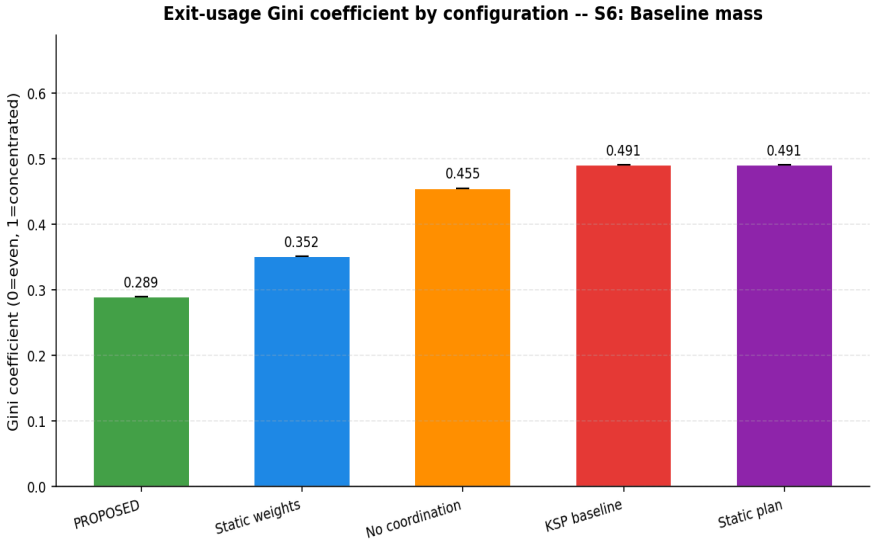


Figure 5. Gini coefficient of exit load, by configuration, for Scenario 6 (values averaged over 20 Monte Carlo replicates).

An important trade-off is also observed. In this group, PROPOSED shows a higher, not lower, mean evacuation time (27.17 s versus 24.98 s) relative to the uncoordinated configurations, despite a significantly better distribution across exits. This is an expected trade-off between speed and load fairness. Coordination deliberately diverts some agents from the nearest, but potentially overloaded, exit toward a more distant but less occupied one, which increases the individual path length for some agents in exchange for a lower risk of congestion at the most heavily used exit.

4.5.4. Group C: IFC-Derived Connectivity, Large Buildings (Scenarios 9–13)

The results for Group C are summarized in Table 14.

Table 14. Significance and mean effect size for Group C.

Metric	vs. STATIC_WEIGHTS	vs. NO_AGENTS_COORD	vs. KSP_BASELINE	vs. STATIC_PLAN
Mean evac. time	0/5 ($\delta = 0.07$)	3/5 ($\delta = 0.53$)	3/5 ($\delta = 0.58$)	4/5 ($\delta = 0.59$)
Max. evac. time	0/5 ($\delta = 0.01$)	5/5 ($\delta = 0.75$)	4/5 ($\delta = 0.61$)	5/5 ($\delta = 0.75$)
Number of exits	0/5 ($\delta = 0.01$)	3/5 ($\delta = 0.51$)	4/5 ($\delta = 0.60$)	4/5 ($\delta = 0.60$)
Gini coefficient	1/5 ($\delta = 0.14$)	2/5 ($\delta = 0.40$)	2/5 ($\delta = 0.31$)	2/5 ($\delta = 0.40$)

Unlike Group B, in the representative Scenario 10, coordination improves the Gini coefficient (0.231 versus 0.456 for NO_AGENTS_COORD) without paying a cost in higher mean time (120.77 s versus 121.57 s). This difference relative to Group B is due to scale: with 48–120 agents and a more complex topology, the available buffer of alternative routes is large enough to allow load balancing without materially lengthening the individual path of any agent. The distribution of evacuated agents across exits for Scenario 10 specifically, under the PROPOSED configuration, is shown in Figure 6. The plotted data represent direct console output from the simulation script (dynamic_simulation.py) and illustrate, at the level of individual agents, how the coordination mechanism distributes the load among the six available exits of the ZDB building, rather than concentrating it on the nearest one.

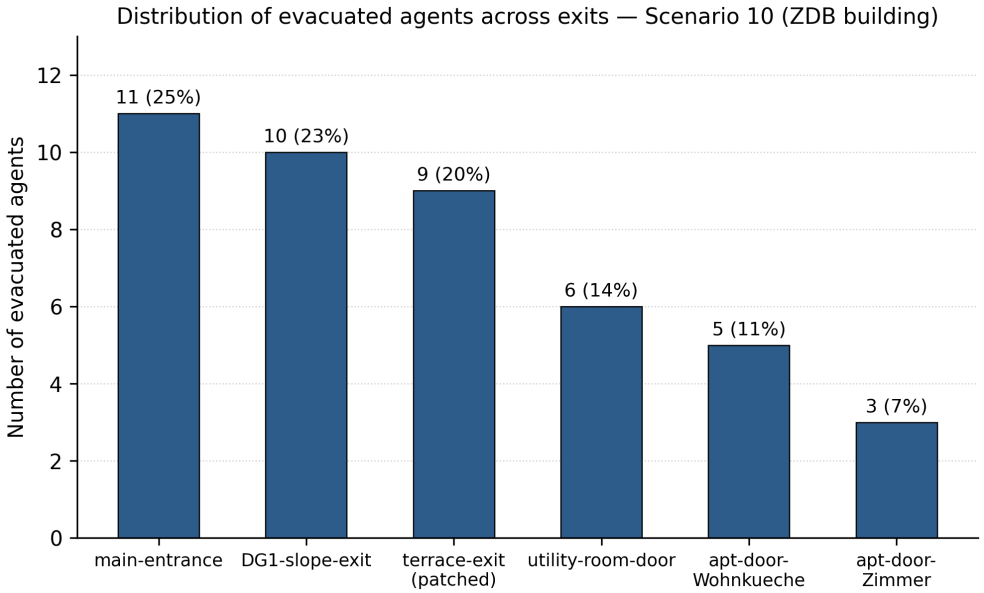


Figure 6. Distribution of evacuated agents across exits for Scenario 10 (ZDB building), extracted from the console output of dynamic_simulation.py for a single illustrative run.

The maximum evacuation time differs significantly from NO_AGENTS_COORD in all five scenarios of this group, but mostly in the direction of a longer worst-case time for the proposed system (Scenarios 10–13, e.g., 75.5 versus 61.0 s in Scenario 11; the number of evacuated agents is equal), and in the opposite direction only in Scenario 9. This is the cost of redistributing the load over exits. The benefit of dynamic replanning is visible only in Scenarios 8 and 9, where the proposed system evacuated more agents.

Scenario 9 is an exception within the group: the building is constructed with a single physically accessible exit, so the Gini coefficient is trivially 0.0 for each of the five configurations, in every one of the twenty replicates, without exception. The Mann–Whitney test on two identical constant samples is not informative. This scenario was deliberately included in the study design to test the algorithm's behavior under a topology devoid of choice. The result confirms that the algorithm does not create artificial variation where no physical choice exists.

4.5.5. Conclusions

The three groups, considered together, trace a consistent picture, in agreement with the architectural analysis from Section 4.4. Calibration of the weight vector shows no measurable effect in any of the three groups—a confirmation of Mechanism B at the level of final simulation results. Coordination and dynamic replanning show an effect whose strength increases monotonically with the size of the sample of possible outcome configurations, as expected from a statistical standpoint. The trade-off between speed and load balance is scale-dependent—noticeable in Group B for smaller, denser buildings, practically absent in Group C at larger scale. The sole structural exception, Scenario 9, is fully explainable by the building's topology.

All conclusions on coordination in this section hold for $\lambda = 0.35$ only; their dependence on λ, building scale, agent population and number of exits were not tested in the main experiments (Section 6.3).

5. Verification of Regulatory Compliance

5.1. Regulatory Requirement

Besides effectiveness, an evacuation algorithm intended for real-world operation must meet a categorical regulatory requirement arising from Regulation No. Iz-1971 [42]: agents must not be routed along a path passing through a node with an active fire, nor through a smoke-affected node. To verify that the proposed architecture satisfies this in practice, a targeted audit of the planner's code was conducted, followed by quantitative verification against the actual simulation logs.

Throughout this article, the "hazard-avoidance property" denotes the following model-level statement: given the sensor classification at a planning time, the recommended route contains no node classified as fire- or smoke-affected. It is not a physical safety guarantee. It does not cover hazards that are not yet classified (missed or delayed detection), agents who are not evacuated within the simulated horizon, or conditions after the planning time.

5.2. Mechanism for Excluding Fire and Smoke Nodes

The planner applies a single, unified principle of hard exclusion of hazardous nodes from the planning graph, applied symmetrically for both smoke and fire. The prohibition on passing through smoke nodes can be disabled programmatically. Smoke is excluded through a tolerance-checking mechanism configured to a strict zero value. This means that any node classified as smoke-affected drops out of the graph at every (re)planning cycle. The same principle applies to all nodes with an active fire: they do not participate in route search as candidates for traversal, regardless of the computed cost of alternative paths. In this way, the selected route contains no classified fire or smoke node at every planning step, by construction and not merely in the statistically more likely case.

5.3. Handling of a Starting Position Within a Hazardous Zone

When an agent is in a node that, at a given moment, is classified as fire- or smoke-affected, the planner activates logic to move the search's starting point to the safest neighboring node. In this case, the abandoned node (the agent's former position) is also excluded from the graph for the purposes of the subsequent search. This ensures within the model that the route cannot subsequently pass back through it at a later step. The combination of the two mechanisms ensures that the recommended route never contains a classified fire or smoke node—neither as an intermediate point nor through a return to an already-vacated hazardous position.

5.4. Verification

Verification was conducted at two levels. An audit of the recommended route, applied across all thirteen scenarios, shows zero instances of a fire or smoke node in the recommended route after the correction, including Scenario 10. An audit of the executed movement—a stricter test, based on the agents_fire and agents_smoke metrics from Section 4.2. The result shows zero incidents across all one hundred runs (five configurations, each with twenty replicas) for twelve of the thirteen scenarios, summarized in Table 15.

Table 15. Number of evacuated agents with physical contact with fire (*agents_fire*) or smoke (*agents_smoke*) during movement, summed over all one hundred runs (5 configurations × 20 replicates) per scenario.

Scenario	Agents_Fire	Agents_Smoke
1–11, 13 (12 scenarios)	0	0
12	230	230

5.5. Boundary Case (Scenario 12)

Scenario 12 is the only exception in Table 15: on average, about two evacuated agents per run show physical contact with fire or smoke during their movement. The number of incidents is virtually identical for all five ablation configurations in each of the twenty replicates—a strong indication that the cause lies outside the scope of the planner's decisions, since if the number of incidents depended on route selection, the better-coordinated configurations would show a lower number.

A direct check of the position log confirmed this hypothesis: in every one of the affected cases, the agent was already located in the affected node now the sensor system classified it as fire- or smoke-affected, since the agent's position before and after classification coincides. The agent did not "enter" the hazard because of a routing decision but was overtaken by it in place. The reason is that the speed of fire spread in this configuration is related to the sensor-check interval, illustrated in Figure 7. In this scenario, the agent physically does not have time to leave the node between the moment of detection and the moment of classification.

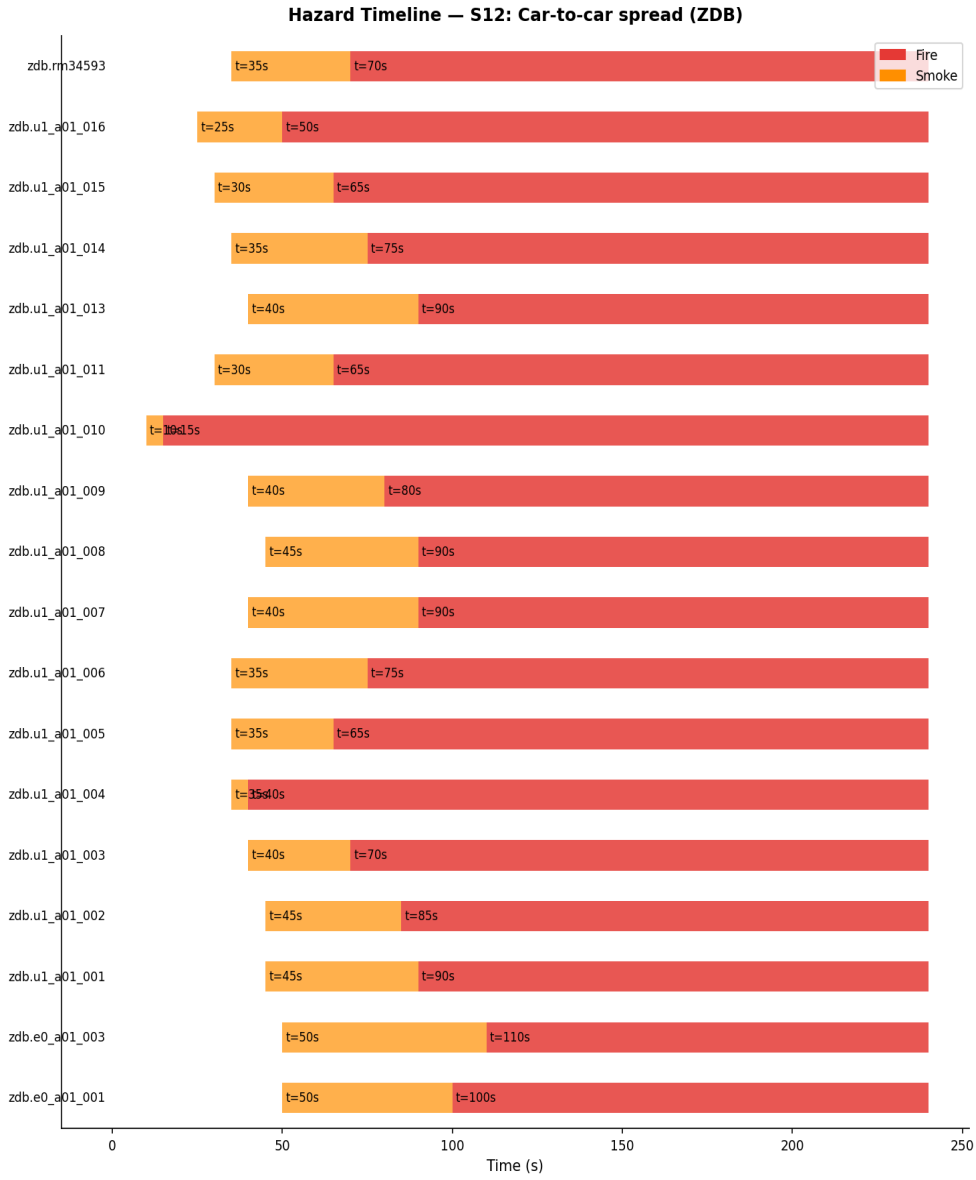


Figure 7. Temporal dynamics of fire and smoke spread across nodes for Scenario 12, showing the rapid transition from smoke (orange) to active fire (red) at each individual node.

Sensor-Interval Sensitivity and Preconditions of the Hazard-Avoidance Property

To quantify the effect of the sensing interval, Scenario 12 was repeated with sensor intervals of 1, 2.5 and 10 s (PROPOSED configuration, 20 paired replicates, movement step fixed at 5 s), and Scenarios 11 and 13 with 1 s (10 replicates each). Contacts with fire or smoke in Scenario 12 fell from 46 across 20 replicates at 5 s to 3 at 1 s, were unchanged at 2.5 s (agents move in 5 s steps) and at 10 s (43). In Scenarios 11 and 13 no contacts occurred at either 1 s or 5 s. Faster sensing was not cost-free: the mean number of evacuated agents decreased in all three scenarios (Scenario 12: 88.9 to 79.1; Scenario 11: 116.0 to 105.6; Scenario 13: 44.1 to 40.3), with a corresponding increase in blocked agents and about a fivefold increase in replanning events. We did not isolate the mechanism of this decrease; it is consistent with earlier classification excluding nodes sooner and cutting off agents who at 5 s still passed through them. The contacts at 5 s therefore depend on the sensing interval relative to the speed of fire spread (they fell only at 1 s) and not on routing decisions, and a shorter interval trades them for blocked agents.

The hazard-avoidance property of Section 5.1 holds under the following preconditions: (i) it refers to nodes classified as hazardous, not to the physical hazard state; (ii) no agent occupies a node at the moment it is classified, which requires the sensing interval to be short relative to the speed of fire spread (Scenario 12 is a case in which it is not); (iii) an admissible route avoiding all classified nodes exists, otherwise the agent is reported as blocked; (iv) route delivery is instantaneous and compliance is complete. The sensing interval and the movement step are design choices, not derived from a standard or from measurements, and the analysis is limited to these scenarios and to a fixed five-second movement step.

5.6. Classification-Based Tenability Margin (Self-Consistency Check)

The verification in Sections 5.1–5.5 establishes a categorical, binary compliance criterion (no hazard node on the recommended or executed route). To complement the binary criterion with a continuous measure, a tenability margin was computed against the onset times of the system's own hazard classification. This is a self-consistency check, not a verification of compliance with ISO 13571 or a classical ASET–RSET comparison.

For each node, ASET_hazard(node) is defined as the earliest simulation timestamp at which the node is classified as fire- or smoke-affected by the sensor-fusion module. For each agent and each simulation step, a spatiotemporal tenability margin is computed as the difference between this onset time and the agent's actual time of occupancy at that node (Equation (12)):

$$\text{margin}(\text{agent},t) = \text{ASET_hazard}(\text{node}(\text{agent}, t)) - t. \quad (12)$$

A negative margin indicates that the agent occupied a node at or after it became hazardous. The worst-case margin per agent is the minimum of this value over its entire trajectory. This measure is a self-consistency check against the system's own hazard model (Section 3.2) rather than an independent ASET derived through the formal Fractional Effective Dose methodology of ISO 13571 [32] from raw, unclassified sensor values; the latter is identified as future work in Section 6.6. ASET_hazard here is the classification onset time, and RSET_max in Table 16 is the maximum simulated evacuation time without a safety factor.

The analysis was applied to the reference simulation run for each of the thirteen scenarios. Table 16 reports, for evacuated agents, the maximum RSET and the worst-case tenability margin, and, separately, the number of agents found "at risk" (negative margin) among evacuated and among non-evacuated agents.

Table 16. Classification-based tenability margin per scenario, evacuated vs. non-evacuated agents.

Scenario	RSET_Max (s)	Worst Margin, Evacuated (s)	At-Risk Evacuated	At-Risk Non-Evacuated
1	40.0	20	0	0
2	50.0	25	0	0
3	80.0	10	0	0
4	55.0	n/a *	0	0
5	30.0	45	0	0
6	45.0	n/a *	0	0
7	45.0	n/a *	0	0
8	60.0	15	0	0
9	40.0	5	0	13
10	275.0	10	0	3
11	85.0	5	0	0
12	85.0	−10	7	23
13	55.0	5	0	27

* "n/a": no node visited by any agent was classified hazardous within the simulated horizon in that run.

Worst margin, evacuated (s) is the minimum spatiotemporal margin (Equation (12)) recorded over the trajectories of evacuated agents; At-risk evacuated/At-risk non-evacuated are the counts of evacuated and non-evacuated agents, respectively, for which this margin was negative at least once.

Two findings emerge. First, for evacuated agents, the margin is positive in every scenario except Scenario 12, exactly reproducing the boundary case already identified and explained in Section 5.5. This is a continuous restatement of the binary result of Table 15 on the same hazard classification, not an independent confirmation. Second, and not previously visible in Table 15, a non-zero number of non-evacuated agents in Scenarios 9, 10, 12, and 13 occupy a node that is subsequently classified as hazardous. Since the *agents_fire*/*agents_smoke* metric in Section 4.2 and Table 15 is computed only over the executed paths of evacuated agents, it does not capture this exposure. This is disclosed here as a genuine limitation of the regulatory-compliance verification as originally scoped, rather than as a newly introduced deficiency of the algorithm: agents that fail to evacuate within the simulated horizon (Section 4.2, definition of the evacuation-rate metric) were, for the same reason, already excluded from the safety-relevant statistics reported in Section 4. Extending the hazard-avoidance property itself (Sections 5.2 and 5.3) to non-evacuated agents, and reporting their tenability status explicitly, is added as a priority item to the future-work discussion in Section 6.6. In the reference runs of Scenarios 9–13, every non-evacuated agent (41, 4, 3, 23 and 27 agents, respectively) received a rescue notification, issued 5–75 s after the start, caused by structural isolation in 3 agents each in Scenarios 9 and 11 and by fire dynamics in all others. Hence this population is not left without a response (a rescue notification), but the metrics of Table 15 do not cover its physical exposure, which is reported only through the margins of Table 16.

5.7. Hydraulic Flow-Capacity Plausibility Check

Sections 4 and 5.1–5.6 validate the proposed system internally, either by comparing it against ablated variants of itself or by auditing its own routing decisions against its own

hazard classification. Neither approach can establish whether the simulated evacuation-completion times are consistent with an evacuation-time estimate obtained by a method entirely independent of the proposed architecture. To close this gap without requiring licensed access to a commercial simulator, this section cross-checks the simulated results against the hydraulic flow model of Nelson and MacLennan [43], the established fire-engineering method that also underlies the movement-time calculations used internally by Pathfinder (in its SFPE calculation mode) and by buildingEXODUS.

For an exit of clear width W, the effective width is:

$$W_e = W - 2b, \tag{13a}$$

where $b = 0.15$ m is the standard boundary-layer allowance per side [43]. The minimum time required for N occupants to pass through that exit is then:

$$t_{flow} = N/(F_s W_e), \tag{13b}$$

where $F_s = 1.3$ persons/(s·m) is the specific flow rate for doors [43]. This is a physical lower bound: no evacuation can complete faster than the time needed to pass the full occupant load through its busiest exit, regardless of routing strategy, sensor behavior, or coordination logic.

For each of the five highest-occupancy, congestion-dominated scenarios (Group C: Scenarios 9–13), W was taken directly from the IFC-derived door geometry of the corresponding building (Institute for Scenario 9; ZDB for Scenarios 10–13), and N was taken from the actual number of agents recorded, per exit, in the simulation log for the PROPOSED configuration. Equation (13b) was evaluated for every exit used in that scenario; the busiest exit, i.e., the one with the largest t_{flow}, gives the binding physical constraint for that scenario. Table 17 reports this bottleneck value alongside the simulated maximum evacuation time already reported in Section 5.6.

Table 17. Hydraulic flow-capacity cross-check against the Nelson–MacLennan model, Group C scenarios.

Scenario	Evacuated Agents	Simulated Max. Evacuation Time (s)	Bottleneck Exit Load (Agents)	Hydraulic t_{flow} (s)	Ratio (Simulated/Hydraulic)
9	39	40.0	39	25.0	1.6×
10	44	275.0	9	7.7	35.8×
11	117	85.0	30	25.6	3.3×
12	97	85.0	31	26.5	3.2×
13	53	55.0	17	10.9	5.0×

In every scenario tested, the simulated maximum evacuation time exceeds the hydraulic lower bound. This establishes only that the simulated times satisfy a necessary physical constraint (physical plausibility); it does not demonstrate quantitative agreement between the pedestrian-flow sub-model and hydraulic theory, which is not claimed here. The ratios in Table 17 should not be read as a measure of model error, for three reasons. First, the simulated time is a clock time that includes pre-movement delays (e.g., in Scenario 10 the last agent starts moving at 245 s and needs about 30 s to exit), whereas t_{flow} contains none. Second, the simulated value refers to the last agent leaving the whole building, while t_{flow} refers to the busiest exit only. Third, the hydraulic bound ignores the travel time to the exit, which is substantial in the multi-storey ZDB building (up to 12 route steps), and the discrete five-second step and density-dependent speed (Equation (9)) add further delay. The ratio is therefore smallest in Scenario 9, which has a single exit, no pre-movement

delay and congestion at one door (1.6×), and largest in Scenario 10 (35.8×), where pre-movement delays dominate. A direct check of the discharge rate at the single exit of Scenario 9 (20 replicates; effective width of 1.2 m derived from Table 17) gives a mean throughput of about 1.2 agents/s, i.e., 76% of the hydraulic capacity of 1.56 agents/s, with the evacuation never faster than the hydraulic bound. Within single five-second steps up to 15 agents (mean peak 10.3) leave, more than the hydraulic capacity of 7.8 per step, because capacity is applied per node and not per door width. The ratios in Scenarios 11–13 reflect mainly travel distance and route complexity (up to 12 route steps) rather than door capacity, while in Scenario 10 they are dominated by pre-movement delays.

This cross-check is an order-of-magnitude plausibility check and is deliberately limited in scope: it does not establish that the system reproduces real-world evacuation timing, since it verifies a physical capacity constraint rather than observed drill data. A full scenario-by-scenario replication in an established commercial simulator such as Pathfinder or buildingEXODUS, discussed as a remaining limitation in Section 6.6, would be required for that stronger claim. What this section establishes is an independent plausibility check, using a method external to the proposed architecture. A quantitative comparison would require the hydraulic time to be augmented with pre-movement and travel time and the simulated discharge rate at each exit to be compared with $F_s W_e$; this is left for future work.

5.8. Tool for Interactive 3D Auditing of a Selected Route

In addition to the quantitative verification in the preceding subsections, the system includes a tool for interactive 3D reconstruction of the route of an arbitrarily selected agent in an arbitrarily selected scenario (the evacuation_3d_visualizer.py module, which uses building geometry extracted from the IFC model via the ifc_geometry_extractor.py module). An example of the evacuation path of agent Z001 in Scenario 11 (ZDB building) is shown in Figure 8.

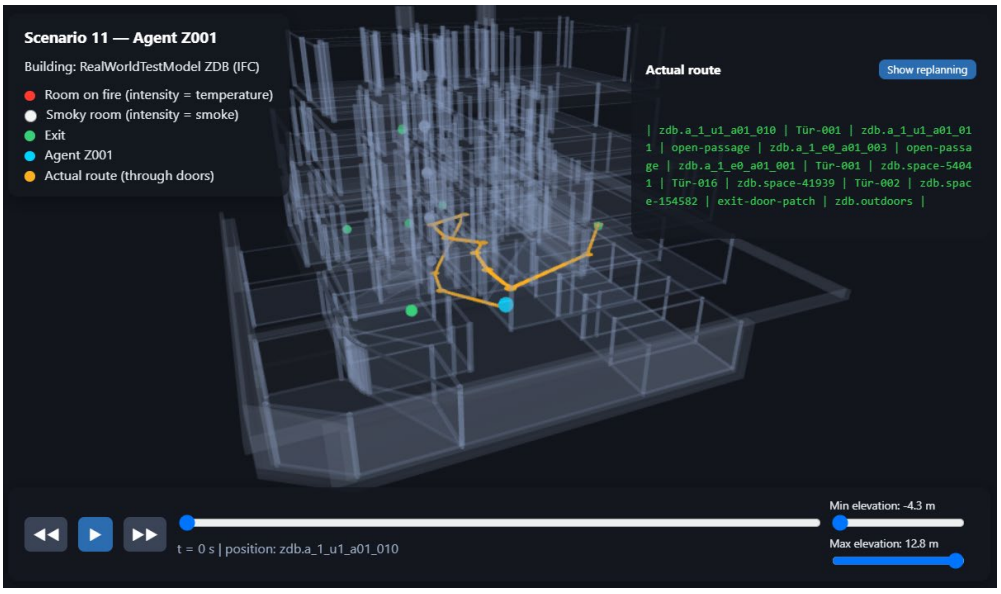


Figure 8. Interactive 3D reconstruction of the actual route of agent Z001 in Scenario 11 (ZDB building), with fire and smoke nodes marked, the actual sequence of doors and spaces traversed, and a slider for time animation and elevation range.

The tool animates the agent's movement through the building space, as well as the development of fire and smoke in the rooms. The higher the temperature or the higher the smoke concentration, the redder or whiter the color of the rooms.

The tool shows not only the final selected path but also the sequence of specific spaces and doors through which the agent passes (the "Actual route" panel), as well as an option to visualize the moments of route replanning ("Show replanning")—a direct visual link to dynamic recomputation. It is available only for scenarios involving buildings with an available IFC model (Scenarios 9–13), since it requires a geometric envelope extracted from the BIM data.

5.9. Conclusions

Sections 5.1–5.5 establish, and Section 5.6 restates through a continuous classification-based margin, that the route recommended by the planner never contains a fire or smoke node for evacuated agents, with the sole physically explained exception of Scenario 12. Section 5.6 additionally identifies non-evacuated agents as a population not covered by this hazard-avoidance property; each such agent triggered a rescue notification in the reference runs, which is carried forward as a limitation in Section 6.6.

It has been verified across all thirteen scenarios that the route recommended by the planner never contains a fire or smoke node. The sole class of recorded incidents involving physical contact between an agent and fire or smoke has a purely physical origin. It arises when the agent is already located in a node at the moment it is classified as hazardous, rather than because of a routing decision.

6. Discussion

6.1. Analysis of the Results

The three independent lines of analysis converge into a single, consistent explanation rather than three separate observations: the sensitivity analysis (Section 4.4), the comparative Monte Carlo study (Section 4.5), and the regulatory-compliance verification (Section 5). The reason all eight of the eight calibratable weights show limited sensitivity is the same reason weight-vector calibration shows no statistically measurable contribution to any output metric in the comparative study: the lexicographic criterion $sort_key = (risk_level, total_cost)$ (Equation (6a)) ensures that the calibratable cost decides only between routes with equal *risk_level*. This architectural feature is also the premise on which the regulatory-compliance mechanism from Section 5 is built. The hard exclusion of fire and smoke nodes from the planning graph operates independently of any weight configuration, which explains why the verification in Section 5.4 finds no violation of the property in any of the five ablation configurations.

The real, statistically confirmed contribution of the proposed algorithm stems from two mechanisms distinct from weight calibration: coordination among agents through virtual load, and dynamic route recomputation upon a change in hazard. Both mechanisms bypass the lexicographic risk-gate, since they do not participate in determining *risk_level*. Coordination affects the distribution of agents among exits that are already equally safe, while dynamic replanning responds to a change in the hazard classification itself, rather than to a change in the calibrated cost.

6.2. Architectural Conclusions

The finding from Section 4.4 has significance beyond the specific implementation examined in this paper. Any evacuation-route-selection architecture that applies a hard, priority safety criterion (as *risk_level* does in the present system) ahead of a calibratable cost component structurally limits the extent to which that component can influence the final decision, regardless of how precisely it is calibrated. This is not a shortcoming. On the contrary, a system in which a calibratable component could override a risk assessment would represent a more substantial safety risk, since an error or drift in calibration could directly

lead to a more dangerous choice. The architectural decision applied in the present system follows a safety-first principle, familiar from other areas of critical-systems engineering: the learned or calibratable components are permitted to optimize only within the bounds already set by a deterministic safety criterion.

The practical implication for researchers working with similar systems is methodological: demonstrating that a calibrated or "trained" component of a given system has no measurable effect on the result is not necessarily an indication of a calibration error or of unsuitable training data. It can be a direct, verifiable consequence of the architecture in which the component is embedded, as shown empirically in Section 4.4.3 through direct instrumentation of the algorithm rather than merely observation of the aggregated results.

6.3. Coordination Design

The results from Section 4.5 reveal a trade-off between evacuation speed and load uniformity across exits, whose direction and magnitude depend on building scale. In smaller buildings (Group B, Section 4.5.3), coordination substantially improves load balance (Cliff's delta up to 1.00) at the cost of a higher mean evacuation time. Agents are deliberately redirected from a nearby but potentially overloaded exit toward a more distant but less occupied one. In larger buildings with richer topology (Group C, Section 4.5.4), the same mechanism improves load balance at practically no cost in mean time, since the available buffer of alternative routes is large enough to absorb the redistribution without materially lengthening individual paths.

Further analysis of the variance of maximum evacuation time (a metric reflecting the predictability of the worst-case evacuation outcome, not merely its mean duration) shows a lower variance of the proposed system than of the baseline with neither coordination nor calibrated weights in the largest-scale scenarios. In Scenarios 11 and 12, the standard deviation of the maximum evacuation time under PROPOSED is 4.26 s and 4.87 s respectively, versus 11.68 s and 23.68 s for the uncoordinated, statically weighted baseline algorithm (KSP_BASELINE)—a reduction on the order of 3 to 5 times (see Figure 9). This pattern is not universal across all thirteen scenarios (PROPOSED shows the lowest variance in eight of thirteen scenarios), but it appears in the two largest scenarios, where predictability of emergency-response time matters most. This comparison is descriptive; the reduction is not observed relative to the variant without coordination, so it cannot be attributed to coordination alone.

The practical implication for deployment is that the value of the coordination coefficient λ (Equation (8)) probably should not be a fixed universal constant, but a parameter whose calibration accounts for the scale and topology of the specific building. This is a conclusion that the present study formulates based on the observed dependence across scenario groups, but does not test directly through a dedicated sensitivity analysis on λ.

All coordination results reported in this paper were obtained with $\lambda = 0.35$ and are valid only for this value. Because the working density is capped at 1.0 (Equation (8)), $\lambda = 0.35$ saturates a node after three planned agents, so λ effectively sets how many agents are assigned to a route before later agents are diverted. This implies a trade-off. If λ is too small, agents concentrate on the nearest exits, as in the uncoordinated case. If λ is too large, routes saturate early, more agents are diverted to distant exits, and once all alternatives are saturated the virtual load no longer discriminates between them, so the balancing effect is lost. The best value is expected to depend on the number of agents, the number of accessible exits and the spare route capacity, consistent with the scale dependence in Sections 4.5.3 and 4.5.4. A limited exploratory sweep ($\lambda \in \{0, 0.1, 0.2, 0.35, 0.5, 0.75, 1.0\}$; Scenarios 6, 8, 10 and 11; five replicates per setting; not part of the main experiments) shows that the dependence of exit-load balance on λ differs between scenarios: in Scenarios 6 and

10 the balance was best near $\lambda = 0.2$ and deteriorated at larger values, whereas in Scenario 8 the Gini coefficient increased with λ, and in no run was there contact with fire or smoke. These runs are too few for statistical conclusions, and a dedicated sensitivity analysis of λ is the priority for future work (Section 6.7).

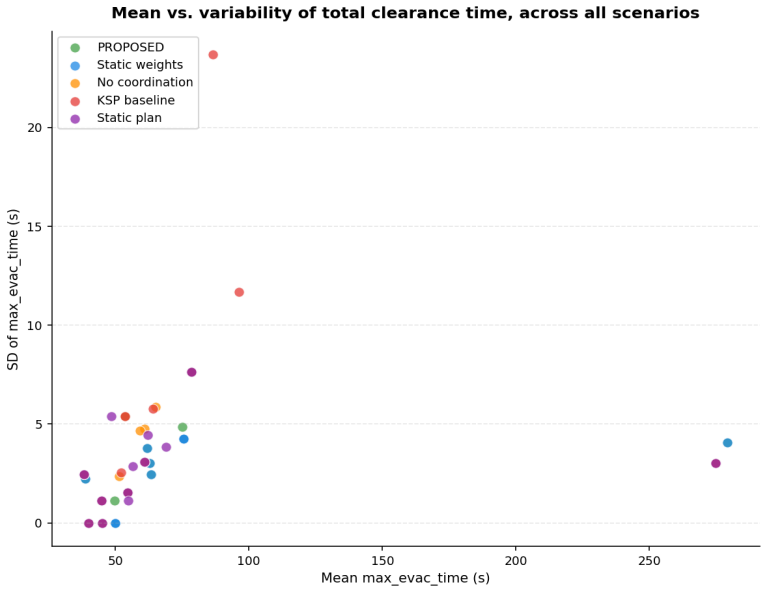


Figure 9. Mean versus standard deviation of maximum evacuation time, by configuration, for all thirteen scenarios with available Monte Carlo results.

6.4. Regulatory Compliance

Section 5 empirically verifies a single, albeit central, regulatory requirement: avoidance of fire and smoke nodes along the route. Beyond this check, the individual mechanisms of the algorithm described in Section 3 are, by design, consistent with a broader set of requirements from the applicable regulatory framework, summarized in Table 18. Unlike Section 5, the correspondences in this table are not subject to separate quantitative verification in the present study and should be regarded as documented design justification rather than as an additional, independently tested result.

Table 18. Correspondence between the algorithm's mechanisms and selected requirements from the applicable regulatory framework (design justification; only the requirement from Section 5 is empirically verified).

Requirement	Regulatory Act	Algorithm Mechanism
Route to a safe zone	Regulation Iz-1971, Art. 31–32	Search restricted to explicitly defined exit nodes (Section 3.3)
Minimum of two independent exits	Regulation Iz-1971, Art. 33	Coordination through virtual load (Section 3.7)
Elevators are not an evacuation route	Regulation Iz-1971, Art. 3; Directive 92/91/EEC	Permanent structural exclusion from the graph (Section 3.3)
Dynamic response to a new hazard	Directive 89/654/EEC	Hysteresis classification and recomputation trigger (Section 3.2)
A safer route takes precedence over a faster one	ISO 23932 [44]; SFPE Handbook	Lexicographic sorting (*risk_level*, *total_cost*) (Section 3.5)
Monitoring of route load	Regulation Iz-1971, Art. 33	Crowd density in risk_score, edge_cost, and node_priority (Sections 3.4 and 3.5)
Notification for blocked individuals	Directive 89/654/EEC; SFPE Handbook	Rescue-notification mechanism (Section 3.7)

6.5. Computational Considerations

The proposed algorithm uses a value of K = 5. A direct comparison with a larger, but practically applicable, value (K = 200) shows that, for twelve of the thirteen scenarios, the number of evacuated agents remains completely identical to that at K = 5, while the runtime of a single execution of the algorithm increases by approximately a factor of 12. In the single scenario with a measurable difference in outcome (Scenario 9), K = 200 yields one fewer evacuated agent, not more. This is a counterintuitive result, but a possible explanation is that, with a significantly larger number of candidate routes sharing the same minimal primary value *risk_level*, the secondary tie-break criterion *total_cost* begins choosing among a much wider set of practically equivalent alternatives, which can lead to different, but not necessarily better, coordination among agents in a topology with a single critical congestion point.

This finding should be read within the conditions under which it was obtained: the lexicographic ranking (*risk_level*, *total_cost*) of Section 3.5, only two values of K (5 and 200), the number of evacuated agents and the algorithm runtime as the compared quantities, thirteen scenarios in four buildings with up to 120 agents, and $\lambda = 0.35$. Under these conditions, additional candidates beyond those covering the different *risk_level* values typically differ only in the secondary term *total_cost*, which has limited influence (Section 4.4.3, Mechanism B), so a larger K incurs mainly computational cost. We do not claim that this holds for other architectures or ranking criteria, e.g., where *total_cost* alone decides, for which more candidates may change the choice, nor for other metrics such as evacuation time and exit-load balance, which were not compared across K. A practical recommendation for systems with similar architecture is that K should be determined empirically, through comparison of the output metrics at increasing values, rather than accepting the value specified in a configuration file without verification.

6.6. Limitations of the Study

Validation through simulation only, and simplified movement behavior. All results in Section 4 stem from a cellular-automaton simulation, not from a real evacuation or drill; absolute evacuation times should not be read as real-world predictions. The comparative conclusions between configurations (Section 4.5) are more robust, since they rely on a paired, controlled comparison under identical hazard conditions, but they remain bound by the assumptions of the movement model. In particular, agent walking speed depends only on local crowd density (Equation (9)) and does not account for smoke-dependent speed reduction [45], physiological effects of smoke inhalation, age/disability differences in mobility, or spatial familiarity; the per-agent mobility attribute exists as a hook for the latter but was not populated in this evaluation. Integrating these established behavioural factors into the movement layer, while preserving the deterministic routing architecture that is this paper's primary contribution, is identified as a priority for future work. The movement layer was not validated against fundamental diagrams or evacuation-drill data; the discharge check of Section 5.7 shows that the coarse five-second step and the node-level capacity allow short bursts above door capacity. The model is therefore suited to comparing routing and coordination strategies under identical conditions and not to predicting absolute evacuation times.

Idealized instruction delivery and compliance. The evaluation in Sections 4 and 5 assumes that every agent instantaneously receives its assigned route and follows it with perfect compliance. Real evacuees do not behave this way. Route delivery depends on a communication channel, such as a smartphone application, a wearable device, or dynamic signage and public-address announcements, that not every occupant will possess, notice, or trust in the moment. Even when a route is received, human decision-making under

fire conditions is bounded rather than optimal: evacuees often default to the most familiar exit instead of the nearest or safest one, delay action while interpreting ambiguous cues, or follow a group or a perceived authority figure rather than an individually assigned instruction. Pre-movement time distributions such as those used in this study [38] account for part of the resulting delay, but not for divergence from the assigned route once movement begins. The results in Sections 4 and 5 should therefore be read as characterizing the routing and coordination logic under idealized compliance, not as a prediction of building-wide outcomes under real human behavior. A probabilistic per-agent compliance model, explicit modelling of familiarity-driven exit choice, and consideration of the instruction-delivery channel's own reliability are identified as priorities for future work.

Latency of sensor detection. Scenario 12 (Section 5.5) is a boundary case. There, the time between the physical start of a hazard and its classification by the sensors can be longer than the time an affected agent needs to leave the node. Section Sensor-Interval Sensitivity and Preconditions of the Hazard-Avoidance Property shows that a 1 s sensing interval removes most of these contacts, but it also reduces the number of evacuated agents. Intervals of 2.5 s and 10 s did not change the contacts materially. We did not isolate the mechanism behind the loss of evacuated agents. The analysis also covers only three scenarios and uses a fixed five-second movement step.

Sensor failure, false alarms and missed alarms. The evaluation assumes complete and correct sensor data, and the present implementation has no fault detection, no staleness check and no redundancy: a missing reading is treated as zero (no hazard), and the last value is retained indefinitely. Three failure modes follow from the architecture. A missed alarm leaves the affected node in the planning graph, so the hazard-avoidance property of Section 5 holds only with respect to nodes that have been classified as hazardous and gives no protection against undetected ones. A false alarm that persists beyond the hysteresis filter removes a node from the graph and may cut off part of the building, producing blocked agents and rescue notifications. A silent sensor failure is indistinguishable from safe reading. Hysteresis suppresses only short transient spikes. We did not test these cases; sensor redundancy and voting, staleness timeouts, fail-safe handling of missing data, and a probabilistic per-agent compliance model are priorities for future work.

Statistical power for small scenarios. Group A (Section 4.5.2) shows an effect size comparable in magnitude to that in the larger groups, but statistically unreliable significance, due to the limited number of possible outcome configurations with 3–8 agents. The conclusions for this group are exploratory and rest on consistency with the larger groups rather than on standalone statistical significance.

Scope of the buildings studied. The thirteen scenarios cover four buildings, represented through two different topology-extraction methods (manual description and automatic extraction from IFC). Generalizability of the findings to buildings with fundamentally different topology (for example, very tall buildings with a central stairwell core, or buildings with irregular, non-hierarchical connectivity) has not been tested directly.

Single regulatory jurisdiction. The conformities documented in Section 6.4 are based on European regulations. The applicability of the same mechanisms to jurisdictions with different requirements (for example, different minimum standards for the number of independent exits, or different smoke-tolerance criteria) requires separate verification against the corresponding local regulations.

Expert-judgement-based fusion and classification parameters. The sensor confidence weights (Table 3) and the risk-score severity coefficients (Equation (2)) were set by engineering judgement rather than by formal calibration or a structured expert-elicitation procedure. The hysteresis thresholds (Table 4) are conceptually consistent with the tenability limits defined in ISO 13571 [32] but were not derived through its formal Fractional Effective Dose

methodology; in particular, the present model treats CO_2 as an independently weighted hazard contributor rather than as a respiration-rate multiplier for other toxicants, as in the FED approach.

We ran a limited exploratory test that was not part of the main experiments. In it, we varied the sensor weights, the entry and exit thresholds, and the normalization maxima by ±20%, using Scenarios 10–12 and five replicates. The exit thresholds had no effect. Mean evacuation time changed by at most about 3%, and the number of evacuated agents by at most about 7%. The number of agents in contact with fire or smoke, however, changed substantially and not in a consistent direction. In Scenario 12, there were 22 contacts in five runs at baseline, 65 with later classification, and 2 with earlier classification. In Scenario 11, there were 0 contacts at baseline and 38 with earlier classification. The hazard-avoidance property itself, which is the hard exclusion of classified nodes, does not depend on these values. What does depend on them is the moment and the extent of classification. In a similar test, the *risk_score* coefficients of Equation (2) had no measurable effect (Section 4.4.5). For these reasons, the system must be calibrated systematically against fire-test data before operational use.

As NFPA 72 [31] confirms, no standard prescribes a specific fusion scheme for combining heterogeneous sensor types into a single hazard score, so this remains a design choice specific to the present implementation. Since these parameters directly affect the primary, safety-governing route-selection criterion (unlike the edge-cost weights in Table 6, whose limited influence is demonstrated in Section 4.4), formal alignment with the ISO 13571 FED methodology, and/or calibration against full-scale or real fire-test data, is identified as a priority for future work.

Non-evacuated agents outside the hazard-avoidance property. The tenability-margin analysis (Section 5.6) shows that agents who did not evacuate are not covered by the hazard-avoidance property. This property was verified in Sections 5.2–5.4, but only for the executed paths of evacuated agents. In the reference runs, every non-evacuated agent received a rescue notification. However, the contact metrics in Table 15 do not show how much these agents were physically exposed to fire or smoke.

Coordination and proximity to unclassified hazards. Classified fire- and smoke-affected nodes are hard-excluded from the routing graph regardless of coordination load, so no exit-balancing detour can route an agent through an already-classified hazard (Sections 3.7 and 5.2). Proximity to a growing but not-yet-classified fire front, however, enters only the secondary tie-break cost term, whose limited influence is shown in Section 4.4. Direct instrumentation across all thirteen scenarios found 965 tie-break events with no instance of a worse-proximity candidate being chosen, though tied candidates tended to share near-identical fire-proximity rather than this being an architectural guarantee. A candidate hardening, extending hard exclusion to a 1-hop buffer around active fire nodes, was tested and altered outcomes in 4 of 13 scenarios (up to a 67% increase in maximum evacuation time in one case), showing this fix is not cost-free. Elevating fire-proximity into the primary risk criterion, or using a smaller, conditional buffer, is left for future work. This limitation is discussed in more detail, together with recommended constraint mechanisms, in Section 4.4.5.

The experiments in Sections 4 and 5.1–5.6 compare the proposed system only with ablated variants of itself and with its own hazard classification. The only independent check is in Section 5.7. It is an order-of-magnitude plausibility test against the Nelson-MacLennan hydraulic flow model. The advantages we report are therefore relative to these ablated variants. We did not compare the system with other published routing methods or with established commercial simulators. As documented by their developers, neither Pathfinder nor buildingEXODUS implements, or can be substituted with, the sensor-fusion

hazard classification, hard node-exclusion, or learned edge-cost mechanisms evaluated here: Pathfinder's fire response relies on manually derived local walking-speed reductions rather than completed FDS coupling [46], while buildingEXODUS depends on an externally supplied hazard field and a Fractional-Effective-Dose-based behaviour model [47]. A scenario-by-scenario replication in either tool, using FDS-generated fire data for a shared building geometry, remains a priority for future work, to establish comparative overall performance under human-behaviour models that this study does not implement.

6.7. Directions for Future Research

Whether the lexicographic route-sorting criterion (TWO_TIER) should be reconsidered in favor of an alternative formulation remains an open architectural question, as discussed in the study's Supplementary Material. Such a change directly affects a decision related to human safety and should be the subject of a dedicated development effort involving a certifying fire-safety engineer, rather than a unilateral algorithmic decision.

The dependence of the trade-off between speed and load fairness on building scale (Section 6.3) motivates a dedicated sensitivity analysis on the coordination coefficient λ, analogous to the one conducted for the eight calibratable weights, with the aim of deriving a recommendation for its value as a function of the number of agents and the number of available exits, rather than a single fixed constant.

The effect of K should be re-examined under other ranking criteria (e.g., *total_cost* alone) and for other metrics such as evacuation time and exit-load balance, since the equivalence of $K = 5$ and $K = 200$ reported in Section 6.5 was obtained only for the lexicographic (*risk_level*, *total_cost*) ranking.

7. Conclusions

The present paper introduced a fully deterministic approach to intelligent fire-evacuation support in multi-storey buildings. The system combines continuous multi-sensor risk assessment with hysteresis, a calibrated composite model for graph-edge costs, a threshold-based table for adaptive node priority, multi-agent coordination through virtual load, and a cellular automaton for modeling evacuee movement. Not a single trainable or "black-box" component is used, with the aim of full traceability, auditability, and compliance with fire-safety regulatory requirements.

The evaluation, conducted across thirteen scenarios in four real buildings (the Rectorate, the University Library, the Institute, and the ZDB building) with up to 120 agents per scenario, shows that offline calibration of the weight vector makes no measurable contribution when coordination is active, a property architecturally explained by the lexicographic risk-gate criterion, which subordinates the calibratable cost to a hard safety priority. The robust contribution of coordination lies in the distribution of load over exits: after correction for multiple comparisons, more exits are used and the Gini coefficient is lower in most scenarios, though not in all (for example, it was higher in Scenarios 4 and 8). This comes with a longer maximum evacuation time in most scenarios where the difference was significant (Scenarios 6, 7 and 10–13), whereas dynamic replanning improved the results only in Scenarios 8 and 9. In the two largest scenarios (11 and 12, 120 agents each), the standard deviation of the maximum evacuation time was 2.7 to 4.9 times lower than for the baseline with neither coordination nor calibrated weights. This descriptive comparison does not hold across all scenarios (the full system had the lowest variance in only eight of thirteen) and was not observed relative to the variant without coordination, so it cannot be attributed to coordination alone. All coordination results correspond to a fixed coefficient $\lambda = 0.35$; their dependence on λ was not tested in the main experiments, and a limited exploratory sweep showed that it differs between scenarios (Section 6.3).

Verification shows that, within the model, the recommended route never contained a classified fire or smoke node. This was held in all thirteen scenarios, for every agent who had an admissible route. The hazard-avoidance property depends on sensor classification, and it does not cover agents who do not evacuate. In the reference runs of Scenarios 9–13, each of these agents received a rescue notification. Their physical exposure is reported only through the tenability margins in Section 5.6.

The only recorded contacts between evacuated agents and fire or smoke happened in Scenario 12. They occurred when an agent was already inside a node at the moment the sensors classified it as hazardous. A shorter sensing interval removed most of these contacts, but it also reduced the number of evacuated agents (Section Sensor-Interval Sensitivity and Preconditions of the Hazard-Avoidance Property). The tenability-margin analysis (Section 5.6) is computed against the system's own hazard classification, so it only restates this result.

Within the simulated scenarios and relative to ablated variants of the system, these results indicate that adaptivity to dynamically evolving fire scenarios can be achieved without trainable or hard-to-verify components. This architectural choice allows full traceability of every evacuation decision, which is a prerequisite for, but not sufficient for, certification of such systems in safety-critical environments; benchmarking against established simulators and published methods remains future work.

Future work will be directed toward integration with high-fidelity CFD models for fire and smoke spread, the use of real-time real IoT sensor streams, extending the system to large-scale, multi-building evacuation scenarios, a dedicated sensitivity analysis of the coordination coefficient relative to building scale, and extending the quantitative verification of regulatory compliance to the full set of requirements documented in the regulatory-positioning table. This study also identifies several further, evidence-grounded directions detailed in Section 6.6: benchmarking against established commercial evacuation simulators such as Pathfinder and buildingEXODUS using FDS-generated fire data; formal calibration of the sensor-fusion and hazard-classification parameters against fire-test data or a structured expert-elicitation procedure, aligned with the ISO 13571 Fractional Effective Dose methodology; extending the hazard-avoidance property to agents who fail to evacuate; elevating anticipatory fire-proximity into the primary safety criterion; fault-tolerant sensing (sensor failure, false and missed alarms); integrating smoke-dependent walking speed, toxicological incapacitation, age- and disability-related mobility, and route familiarity into the movement model; and modelling the reliability of instruction delivery and the bounded rationality of occupant compliance.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/app162111080/s1, software (software folder): contains the Python code used to generate all results (23 Python files); results (folder with all results): contains output files in JSON format with results generated by the software; docs (documents folder): contains three PDF files with project structure, software description, and 3D visualizer description.

Funding: This research received no external funding.

Data Availability Statement: The source code, configuration files, sensor input data, and simulation outputs supporting the findings of this study are openly available as Supplementary Material accompanying this article.

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

References

1. Ivanov, R. An Algorithm for On-the-Fly K Shortest Paths Finding in Multi-Storey Buildings Using a Hierarchical Topology Model. *Int. J. Geogr. Inf. Sci.* **2018**, *32*, 2362–2385. [CrossRef]

2. Qi, X.; Li, X.; Zhang, J.; Song, W. Pedestrian Evacuation in Fire Environments: Scientometric Analysis and Systematic Evaluation. *Phys. A* **2026**, *686*, 131323. [CrossRef]
3. Muthukumar, S.; Rajesh, S.; Talpallikar, A.; Mishra, T.K. A Systematic Review of Artificial Intelligence Frameworks for Holistic Disaster Management. *Discov. Artif. Intell.* **2026**, *6*, 308. [CrossRef]
4. Leong, C.S.; Lau, S.H.; Liew, H.H. Conceptual Fire Risk Management Framework of Building Information Modeling and Fire Dynamic Simulator. *Eng. Proc.* **2025**, *91*, 11. [CrossRef]
5. Zhang, Q.; Hu, Y.; Li, X. Research on Fire Scenario Analysis and Emergency Response Strategies for L-Shaped Buildings Using FDS. *PLoS ONE* **2026**, *21*, e0346927. [CrossRef]
6. Sharma, J.; Andersen, P.A.; Granmo, O.C.; Goodwin, M. Deep Q-Learning with Q-Matrix Transfer Learning for Novel Fire Evacuation Environment. *IEEE Trans. Syst. Man Cybern. Syst.* **2021**, *51*, 7363–7381. [CrossRef]
7. Feng, Z.; Wang, C.; An, J.; Zhang, X.; Liu, X.; Ji, X.; Kang, L.; Quan, W. Emergency Fire Escape Path Planning Model Based on Improved DDPG Algorithm. *J. Build. Eng.* **2024**, *95*, 110090. [CrossRef]
8. Fang, H.; Tam, W.C.; Lovreglio, R.; Emon, M.I.S.; Huang, M.X. Development of a Tenability-Based Path Planning Model for Building Fire Evacuations Using Reinforcement Learning. *J. Build. Eng.* **2026**, *118*, 115132. [CrossRef]
9. Sun, L.; Liu, H.; Fan, B.; Li, X.; Li, W. Multi-Agent Deep Reinforcement Learning-Based Path Planning for Crowd Evacuation in Buildings. *J. Build. Eng.* **2025**, *111*, 113345. [CrossRef]
10. Xie, C.Z.T.; Chen, Q.; Zhu, B.; Lee, E.W.M.; Tang, T.Q.; Yin, X.; Yuan, Z.; Zhang, B. Coordinating Dynamic Signage for Evacuation Guidance: A Multi-Agent Reinforcement Learning Approach Integrating Mesoscopic Crowd Modeling and Fire Propagation. *Chaos Solitons Fractals* **2025**, *194*, 116246. [CrossRef]
11. Omidbeyk, S.; Ghavami, S.M. Multi-Agent Simulation of Population Evacuation during Dynamic Fire Using Reinforcement Learning Based on Integration of Geographic Information Systems and Building Information Modeling. *J. Build. Eng.* **2025**, *111*, 113035. [CrossRef]
12. Deng, H.; Fan, Z.; Wei, X.; Deng, Y. An Ontology-Based Approach to Dynamic Indoor Fire Emergency Evacuation Path Planning with BIM Integration. *J. Build. Eng.* **2025**, *106*, 112562. [CrossRef]
13. Jiang, S.; Wei, Y.; Zhang, B. Fire Evacuation Design Optimization of Public Buildings Based on Building Information Modeling Semantic Enrichment for Fire Simulation Requirements. *Reliab. Eng. Syst. Saf.* **2026**, *277*, 113173. [CrossRef]
14. Shih, G.R.; Tsai, P.H. Safest-Path Planning Approach for Indoor Fire Evacuation. *Int. J. Disaster Risk Reduct.* **2023**, *93*, 103760. [CrossRef]
15. Guan, W.; Hou, S.; Yu, G.; Gong, H.; Guan, S.; Zhao, J. Dynamic Evacuation Path Planning for Multi-Exit Building Fire: Bi-Objective Model and Algorithm. *Fire Technol.* **2023**, *59*, 2853–2876. [CrossRef]
16. Yen, H.H.; Lin, C.H.; Tsao, H.W. Novel Smoke-Aware Individual Evacuation and Congestion-Aware Group Evacuation Algorithms in IoT-Enabled Multi-Story Multi-Exit Buildings. *IEEE Access* **2022**, *10*, 119402–119418. [CrossRef]
17. Li, N.; Huang, G.; Jiang, H.; Gao, X.; Zhou, L. Fire Propagation-Driven Dynamic Intelligent Evacuation Model in Multifloor Hybrid Buildings. *Adv. Eng. Inform.* **2023**, *57*, 102097. [CrossRef]
18. Bai, J.; Lv, X.; Nie, L.; Fang, M. Evacuation Route Determination in Indoor Architectural Environments Based on Dynamic Fire Risk Assessment. *Buildings* **2025**, *15*, 1715. [CrossRef]
19. Xu, L.; Huang, K.; Liu, J.; Li, D.; Chen, Y.F. Intelligent Planning of Fire Evacuation Routes Using an Improved Ant Colony Optimization Algorithm. *J. Build. Eng.* **2022**, *61*, 105208. [CrossRef]
20. Zhang, Y.Q.; Wang, J.H.; Wang, Y.; Jia, Z.C.; Sun, Q.; Pei, Q.Y.; Wu, D. Intelligent Planning of Fire Evacuation Routes in Buildings Based on Improved Adaptive Ant Colony Algorithm. *Comput. Ind. Eng.* **2024**, *194*, 110335. [CrossRef]
21. Zuo, S.; Mao, Z.; Fan, C.; Chen, X.; Gong, M.; Ren, J.; Fan, X.; Guo, Y. Dynamic Planning of Crowd Evacuation Path for Metro Station Based on Dynamic Avoid Smoke A-Star Algorithm. *Tunn. Undergr. Space Technol.* **2024**, *154*, 106145. [CrossRef]
22. Golshani, F.; Fang, L. A Fire Navigation Model: Considering Travel Time, Impact of Fire, and Congestion Severity. *Reliab. Eng. Syst. Saf.* **2025**, *259*, 110914. [CrossRef]
23. Shi, J.; Li, M.; Jiang, N.; Yu, H.; Yang, H.; Zhou, X.; Yang, L. A Multi-Agent Deep Reinforcement Learning Model for Crowd Coordinated Evacuation Simulation in Complex Environments. *Simul. Model. Pract. Theory* **2026**, *147*, 103238. [CrossRef]
24. Yu, A.; Bu, H.; Luan, T.; Gai, W. Integrated Multi-Agent-Based Outpatient Building Fire Response Modeling for Risk-Driven Resource Use and Retrofitting Strategies: A Case Study. *Reliab. Eng. Syst. Saf.* **2025**, *260*, 110970. [CrossRef]
25. Bahmani, H.; Ao, Y.; Yang, D.; Xu, Q.; Zhao, J. Enhancing Evacuation Safety in Urban Primary Schools: An Agent-Based Model Integrating Child Development Behaviour and Health Dynamics. *Reliab. Eng. Syst. Saf.* **2026**, *265*, 111591. [CrossRef]
26. Saravanan, T.; Singhal, U.; Magadum, N.; Ramasamy, P.; Paramanandam, K. AI-Driven Crowd Behavior Simulation for Emergency Evacuation Management. In *International Conference on Data Science and Applications*; Springer Nature: Cham, Switzerland, 2025; pp. 65–78. [CrossRef]
27. Yang, S.; Ceferino, L.; Zhang, Y.; Gu, C.; Guo, T.; Kondo, G. When Agents Learn to Think: Large Language Model-Enhanced Agent-Based Modeling for Crowd Evacuation in Disaster Scenarios. *Reliab. Eng. Syst. Saf.* **2026**, *269*, 112056. [CrossRef]

28. Lin, J.R.; Chen, K.Y.; Song, S.Y.; Cai, Y.H.; Pan, P.; Deng, Y.C. Digital Twin of Buildings and Occupants for Emergency Evacuation: Framework, Technologies, Applications and Trends. *Adv. Eng. Inform.* **2025**, *66*, 103419. [CrossRef]
29. Pishahang, M.; Ruiz-Tagle, A.; Ramos, M.A.; Droguett, E.L.; Mosleh, A. A Bayesian Agent-Based Model and Software for Wildfire Safe Evacuation Planning and Management. *Proc. Inst. Mech. Eng. Part O J. Risk Reliab.* **2025**, *239*, 515–534. [CrossRef]
30. Sun, H.; Chen, M.; Zhang, Y.; Zhang, T.; Hu, P.; Ye, X.; He, Q.; Zhu, K.; Wu, K. EvacuNet: An Intelligent Decision Support Framework for Real-Time Evacuation Assessment in Complex Buildings. *Expert Syst. Appl.* **2026**, *329*, 132898. [CrossRef]
31. National Fire Protection Association. *NFPA 72: National Fire Alarm and Signaling Code*; NFPA: Quincy, MA, USA, 2025. Available online: https://www.nfpa.org/product/nfpa-72-national-fire-alarm-and-signaling-code/p0072code (accessed on 17 September 2026).
32. *ISO 13571:2012*; Life-Threatening Components of Fire—Guidelines for the Estimation of Time to Compromised Tenability in Fires. International Organization for Standardization: Geneva, Switzerland, 2012. Available online: https://www.iso.org/standard/56172.html (accessed on 17 September 2026).
33. *NFPA 92*; Standard for Smoke Control Systems. NFPA: Quincy, MA, USA, 2024. Available online: https://www.nfpa.org/codes-and-standards/nfpa-92-standard-development/92 (accessed on 18 August 2026).
34. *EN 1991-1-2:2002*; Eurocode 1—Actions on Structures—Part 1-2: General Actions—Actions on Structures Exposed to Fire. CEN: Brussels, Belgium, 2002.
35. Alpert, R.L. Calculation of Response Time of Ceiling-Mounted Fire Detectors. *Fire Technol.* **1972**, *8*, 181–195. [CrossRef]
36. *BS 9999:2017*; Fire Safety in the Design, Management and Use of Buildings—Code of Practice. British Standards Institution: London, UK, 2017. Available online: https://knowledge.bsigroup.com/products/fire-safety-in-the-design-management-and-use-of-buildings-code-of-practice (accessed on 18 August 2026).
37. National Fire Protection Association. *NFPA 101: Life Safety Code*; NFPA: Quincy, MA, USA, 2024. Available online: https://www.nfpa.org/product/nfpa-101-life-safety-code/p0101code/nfpa-101-life-safety-code-2024/10124 (accessed on 18 August 2026).
38. *PD 7974-6:2019*; Application of Fire Safety Engineering Principles to the Design of Buildings—Part 6: Human Factors: Life Safety Strategies—Occupant Evacuation, Behaviour and Condition. British Standards Institution: London, UK, 2019. Available online: https://knowledge.bsigroup.com/products/application-of-fire-safety-engineering-principles-to-the-design-of-buildings-human-factors-life-safety-strategies-occupant-evacuation-behaviour-and-condition-sub-system-6 (accessed on 18 August 2026).
39. Karlsruhe Institute of Technology; Institute for Automation and Applied Informatics. AC20-Institute-Var-2 [IFC Test Model]; IfcWiki, 2017. Available online: https://www.ifcwiki.org/images/9/98/AC20-Institute-Var-2.ifc (accessed on 18 August 2026).
40. Fischer, S.; Schranz, C.; Urban, H.; Pfeiffer, D.; Flamm, L. *Real-World Test Model for Escape Route Analysis in IFC Format*; Version 2.0.0 [Data Set]; TU Wien: Vienna, Austria, 2024. [CrossRef]
41. *ASTM E1355-23*; Standard Guide for Evaluating the Predictive Capability of Deterministic Fire Models. ASTM International: West Conshohocken, PA, USA, 2023. Available online: https://store.astm.org/e1355-23.html (accessed on 18 August 2026).
42. Ministry of Regional Development and Public Works. Regulation No. Iz-1971 on Construction-Technical Rules and Standards for Ensuring Fire Safety. 2009. Available online: https://www.mrrb.bg/bg/naredba-z-1971-ot-29-oktomvri-2009-g-za-stroitelno-tehnicheski-pravila-i-normi-za-osiguryavane-na-bezopasnost-pri-pojar/ (accessed on 18 August 2026). (In Bulgarian)
43. Nelson, H.E.; MacLennan, H.A. Emergency Movement. In *SFPE Handbook of Fire Protection Engineering*, 2nd ed.; DiNenno, P.J., Ed.; National Fire Protection Association: Quincy, MA, USA, 1995; pp. 3-286–3-295.
44. International Organization for Standardization. *ISO 23932-1:2018*; Fire Safety Engineering—General Principles, Part 1: General. ISO: Geneva, Switzerland, 2018. Available online: https://www.iso.org/standard/63933.html (accessed on 19 September 2026).
45. Fridolf, K.; Nilsson, D.; Frantzich, H.; Ronchi, E.; Arias, S. Walking Speed in Smoke: Representation in Life Safety Verifications. In Proceedings of the 12th International Conference on Performance-Based Codes and Fire Safety Design Methods, Honolulu, HI, USA, 24–26 April 2018.
46. Thunderhead Engineering. *Pathfinder Technical Reference Manual*; Thunderhead Engineering: Manhattan, KS, USA, 2024. Available online: https://www.thunderheadeng.com/docs/2024-1/pathfinder/technical-reference-manual/ (accessed on 19 September 2026).
47. Galea, E.R.; Lawrence, P.J.; Gwynne, S.; Filippidis, L.; Blackshields, D.; Cooney, D. *buildingEXODUS V6.2 Technical Manual and User Guide*; University of Greenwich: London, UK, 2015. Available online: https://gala.gre.ac.uk/15470/ (accessed on 19 September 2026).

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.