Simulation Modelling Practice and Theory 150 (2026) 103294

Contents lists available at ScienceDirect

Simulation Modelling Practice and Theory

journal homepage: www.elsevier.com/locate/simpat

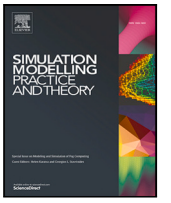

When shortest is not safest: Multi-agent evacuation with awareness and agile routing in dynamic hazards☆

Álvaro Serrano [a], Marin Lujak [a,*], Giuseppe Vizzari [b]

[a] CETINIA, Universidad Rey Juan Carlos, C/ Tulipán s/n, Móstoles, 28933, Spain
[b] Università degli Studi di Milano-Bicocca, Piazza dell'Ateneo Nuovo, 1, Milano, 20126, Italy

ARTICLE INFO

Keywords:
Agent-based simulation
Emergency evacuation
Dynamic risk-aware routing
Situational awareness
Agile evacuation routes
Crowd dynamics

ABSTRACT

Indoor environments can present significant spatial structural complexity, making safe and timely evacuation a critical challenge in the presence of dynamically evolving hazards and incomplete or delayed information available to evacuees. In this context, we study pedestrian evacuation in complex indoor environments using agent-based simulation, focusing on the performance of dynamic routing strategies under varying levels of situational awareness. We consider scenarios in which ambient intelligence systems detect pedestrians and hazards and provide real-time guidance through variable message panels and mobile applications. Using the open-source JuPedSim simulator extended with a customized wayfinding model, we evaluate four guidance configurations that combine low and high situational awareness with shortest-path and agile routing strategies. Low-awareness conditions represent locally informed evacuations, in which pedestrians rely only on reactive, locally available hazard information, while high-awareness scenarios provide dynamic risk updates during the evacuation process. Agile routing prioritizes structurally flexible paths, enabling adaptive rerouting when hazards evolve. Evacuation performance is assessed in terms of evacuation time and remaining-path risk (RPR), computed from the time-dependent node-level risk along the selected remaining route. Simulation results demonstrate that high situational awareness consistently improves both evacuation time and safety. Moreover, agile routing proves particularly effective when hazards emerge late or when spatial layouts are complex, especially under high-awareness guidance that supports adaptive route selection.

1. Introduction

The increasing complexity, size, and density of modern indoor environments, such as public buildings, hospitals, and shopping centers, pose significant challenges for the safe evacuation of large populations. These facilities are often characterized by intricate spatial layouts, limited visibility, and constrained decision-making time during emergencies. The increasing frequency of crowd-related disasters (see the recent report by Feliciani et al. [1]) further highlights the need for effective evacuation strategies that can cope with such conditions. Empirical studies demonstrate that evacuation behavior is strongly influenced by situational awareness, perceived urgency, and the availability and clarity of guidance [2–4]. Together, these findings highlight the need for evacuation models that explicitly consider how information availability and guidance influence dynamic risk-aware routing and evacuation outcomes in evolving hazard conditions [5].

☆ This article is part of a Special issue entitled: 'SIMPAT_ Pedestrian Simulation' published in Simulation Modelling Practice and Theory.

* Corresponding author.
E-mail addresses: alvaro.serrano@urjc.es (Á. Serrano), marin.lujak@urjc.es (M. Lujak), giuseppe.vizzari@unimib.it (G. Vizzari).

https://doi.org/10.1016/j.simpat.2026.103294

Available online 20 May 2026

Historically, wayfinding and guidance mechanisms have proven crucial in improving evacuation outcomes, as shown by studies employing evacuation drills and responses to false alarms [6]. Simulation-based studies have further investigated individual decision-making during routine conditions [7] and in the presence of hazards [8], showing how different wayfinding strategies influence evacuation dynamics when external guidance is absent. More recent research has moved beyond static assumptions by explicitly incorporating situational awareness and adaptive cognition into evacuation models, exploring how dynamic perception, information availability, and communication affect pedestrian responses [9,10].

In parallel to advances in behavioral modeling, research has increasingly questioned the adequacy of shortest-path routing in dynamic emergency scenarios. Pilot studies have demonstrated the feasibility of real-time guidance systems that integrate environmental monitoring and control mechanisms to effectively direct pedestrian movements, particularly in large-scale venues [11,12]. These initiatives provide a foundation for the development of innovative evacuation strategies that consider more than just the shortest-path algorithms. Alternative routing criteria, such as *node centrality* defined over a discrete, graph-based representation of the environment, offer increased flexibility in dynamically evolving scenarios where exits may become unavailable or new hazards emerge. By favoring structurally central nodes, these approaches facilitate the dynamic selection of alternative routes during evacuation, potentially providing a substantial advantage in complex indoor architectures [13,14]. In addition, new hybrid approaches combining physical- and cognitive-based modeling, such as social-force model extensions with awareness mechanisms, are emerging as promising tools to reproduce realistic crowd adaptation in dynamic hazards [15,16].

Assuming the presence of an ambient intelligence infrastructure capable of detecting pedestrians and hazards through devices such as surveillance cameras and smoke detectors, and influencing pedestrian decisions via variable message panels and/or mobile apps, this study investigates the potential impact of different routing strategies under varying levels of pedestrian awareness. We distinguish between low-awareness conditions, in which evacuees rely only on local hazard information and monitor only the next node of their current route, and high-awareness conditions, in which the technological infrastructure provides pedestrians with real-time risk information. In both conditions, evacuees receive route guidance that determines their wayfinding decisions. For the purposes of this study, we assume that pedestrians comply with the provided guidance. This assumption is common in evacuation modeling and particularly when assessing the potential impact of intelligent guidance systems. However, it represents a simplifying limitation, as compliance cannot be guaranteed in real-world scenarios. Empirical studies such as [17] have explored the tendency of pedestrians to comply with guidance recommendations through surveys, while simulation-based analyses (e.g., [18]) have examined the implications of varying levels of compliance (showing that even limited compliance can sometimes bring significant improvements in alleviating congestion). Although explicitly modeling compliance dynamics and trust remains an important research direction, our goal is to estimate the potential performance gains arising from high situational awareness and agile routing strategies.

We employ the open-source simulator JuPedSim [19],[1] extended with a customized wayfinding model and mechanisms for generating dynamic hazards in the environment. Using this framework, we simulate dynamic environments in which risks emerge without pedestrians' prior knowledge, and evaluate four guidance configurations obtained by crossing two awareness regimes (low/reactive vs. high/anticipatory) with two routing strategies (efficient-path vs. centrality-based). We use metrics associated with both evacuation time and path risk, where risk is characterized through statistics of the time-dependent node-level risk along the remaining route.[20]. The achieved results confirm that high situational awareness consistently improves both evacuation time and safety. Furthermore, agile routing offers significant advantages in scenarios involving late-emerging hazards in complex spatial topologies, especially when coupled with high-awareness. These findings highlight the importance of integrating real-time risk information with structurally adaptive routing strategies in evacuation models, with direct implications for the design of intelligent evacuation systems in real-world settings.

Although higher situational awareness intuitively improves evacuation performance, its interaction with the adopted routing strategy remains underexplored. In environments with dynamically evolving hazards, awareness alone does not guarantee safety if the underlying routing strategy lacks structural adaptability. Conversely, even robust routing strategies may fail when agents are unable to perceive and react to evolving risks in a timely manner. Therefore, we focus on the joint effect of situational awareness and routing strategies, aiming to quantify how their interplay influences evacuation efficiency and safety. This combination provides insight into the manner in which real-time information and structural path flexibility jointly influence safety outcomes that cannot be adequately captured when each factor is analyzed in isolation.

The paper is organized as follows: Section 2 reviews the most relevant literature on evacuation modeling, situational awareness, and adaptive pathfinding strategies. Section 3 introduces the routing logic studied in this work, detailing both the efficient and the centrality-aware strategies, together with the situational awareness mechanisms and the associated decision rules. Section 4 then describes the underlying simulation framework, including the modular architecture, the risk modeling component, and the routing and agent modules that implement the proposed behaviors. Next, Section 5 presents the simulation environments and motivates their structural characteristics, explaining how scenario design supports non-trivial route choice and hazard-driven adaptation. Section 6 reports the experimental evaluation, including the setup, performance metrics, and a discussion of results across representative scenarios. Finally, Section 7 summarizes the main findings and outlines directions for future work. Supplementary results and full reproducibility details, including access to scenario configurations and raw outputs, are provided in Appendix A.

[1] https://www.jupedsim.org/

2. Related work

The study of pedestrian behavior during emergency evacuations has evolved significantly over the past two decades, shifting from simple flow models to more sophisticated, microscopic simulations capable of reproducing individual decision-making and emergent group dynamics. A foundational contribution to this field is the social force model introduced by Helbing et al. [21], which describes pedestrian movement as the result of attractive and repulsive forces between agents and obstacles. This model was later extended to simulate panic behavior and crowd phenomena such as herding, clogging, and arching near exits [22]. More recently, Zhang et al. [15] proposed a social-force-based model with static leader guidance, showing how leader placement and influence affect evacuation efficiency in constrained environments.

In parallel with these developments, an active line of research on movement and evacuation dynamics of pedestrians[2] has developed additional microscopic modeling approaches: some of them still interpret pedestrians as particles subject to forces such as the generalized centrifugal force model [23]; others are based on discrete representations of the environment and rules for pedestrian movement following a Cellular Automata approach (see, e.g., the floor field model [24]); other ones adopt autonomous agents approaches, recently also employing Machine Learning approaches, such as neural networks [25] or Deep Reinforcement Learning [26]. Within this direction, the RESCUE framework [16] proposes a high-fidelity, real-time crowd evacuation simulation model aligned with a sensory–decision–motor paradigm, integrating adaptive social-force-based decision mechanisms with personalized motion control to capture realistic individual behavior, interactions, and terrain effects.

To better reflect cognitive processes, Schadschneider et al. [27] adapted a hierarchical behavioral framework that distinguishes between strategic, tactical, and operational decision layers to the pedestrian modeling context. The tactical level, which governs short-term routing and local adaptation, has been emphasized in several studies as critical to effective evacuation modeling [11,28]. In particular, early work demonstrated the feasibility of simulating pedestrian evacuations in near real time for large-scale venues through algorithmic optimization and parallel computation [11]. Numerous agent-based models have since incorporated tactical flexibility. For example, Wagoum et al. [8] studied route changes driven by perceived congestion, while Vizzari et al. [7] introduced cognitively plausible heuristics based on the local context. Andresen et al. [29] considered situations in which pedestrians have different levels of familiarity with the environment in which they are located, and this influences their wayfinding behavior. These works showed that (i) route choice decisions are highly dynamic and consider multiple interacting factors, (ii) they have quantifiable implications for evacuation performance, and (iii) static routing decisions can be insufficient under rapidly changing conditions.

A small but growing number of pilot studies have investigated the plausibility of implementing crowd control centers or dynamic guidance systems that gather environmental information and influence pedestrian movements in real time, particularly in mass egress scenarios such as arenas [11,12]. More generally, additional recent works explored the possibility of influencing crowds through *nudging* techniques (visual, auditory, or acoustic cues) to modulate route choice and dispersal patterns [30,31]. The effect of compliance on guidance effectiveness has also been analyzed both through user studies [17] and simulation-based experiments [18], showing that even partial adherence to guidance can significantly mitigate congestion.

Guo et al. [32] and Lam et al. [20] showed that impaired visibility or limited perception alters evacuation effectiveness, underscoring the need to model agents with varying access to environmental information. Recent studies have deepened this line of research by explicitly modeling perception, communication, and cooperative awareness [9,10]. For instance, Keykhaei et al. [9] simulated human cognition and information sharing in earthquake evacuations, while Zablotsky et al. [10] examined how imitation of cooperative behavior improves collective outcomes.

Recent experimental studies have also highlighted the influence of guidance visibility, stress, and signage design on behavioral responses during evacuation. Zhang et al. [2] demonstrated that guidance cues under varying urgency levels significantly affect exit choice and travel efficiency, while Wang et al. [3] confirmed through virtual-reality experiments that optimized sign height and spacing improve route adherence. Mossberg et al. [4] further evaluated evacuation systems from a situational-awareness perspective, revealing how awareness-based information design can improve crowd coordination.

Real-world incidents have illustrated the critical interplay between human behavior and infrastructure. The analysis of the Love Parade disaster [33] revealed how architectural bottlenecks and communication breakdowns led to tragedy. Similarly, Sieben and Seyfried [34] documented systemic vulnerabilities during mass events, reinforcing the value of simulation as a planning and preparedness tool.

In parallel with advances in behavioral modeling, agile routing has been investigated from both dynamic optimization and network-theoretic perspectives. In this context, an *agile route* denotes the ability to safely and efficiently reroute from intermediate locations when local safety conditions deteriorate. The generation of alternative evacuation routes is commonly grounded in classical formulations of k-shortest loopless paths in networks [35], which provide a systematic way to explore near-optimal routing options beyond purely shortest paths. From a dynamic optimization viewpoint, Lujak et al. [13] proposed a multi-agent architecture for real-time evacuation route selection that accounts for evolving hazards and stress-related pedestrian responses, enabling adaptive guidance through smart environments. From a structural perspective, centrality metrics have been introduced to characterize the importance of nodes within evacuation networks. Building on this idea, Lujak et al. [14] defined evacuation-specific centrality measures and showed that routing strategies favoring structurally central nodes can enhance evacuation efficiency and safety, particularly under dynamically evolving risk conditions.

[2] Also testified by the conference series on Pedestrian and Evacuation Dynamics (PED) - https://collective-dynamics.eu/index.php/cod/ped

From a technological standpoint, simulation platforms such as JuPedSim [19], as well as other open-source or commercial simulators, enable granular and plausible agent-based modeling with modular components for crowd dynamics, hazard propagation, and adaptive behavior. These frameworks support the evaluation of complex evacuation strategies at scale and facilitate the integration of real-time perception modules.

Overall, the literature has made considerable progress in modeling pedestrian evacuation behavior, routing strategies, and situational awareness. The technological landscape now supports adaptive, information-driven guidance and the study of cognitive behavioral responses to evolving risks. However, despite these advances, to the best of our knowledge, the literature lacks a systematic investigation of the *combined effect of agile routing and awareness* on evacuation performance under dynamically changing hazard conditions. This research addresses that gap by providing a comparative analysis of efficient-path and agile routing under both high- and low-awareness conditions, within a unified agent-based simulation framework capable of representing dynamic hazards and awareness-driven agile routing.

3. Routing strategies

This section examines pedestrian evacuation under uncertainty through two routing strategies: (i) *efficient path selection*, which prioritizes nominal travel-time efficiency, and (ii) *agile path selection*, which promotes rerouting flexibility via centrality-based optimization. Both strategies are defined and analyzed within a unified path-selection model that specifies the navigation graph, candidate route generation, and risk-based route adaptation.

3.1. Path selection model

Following [14], we represent the environment as a simple directed navigation graph $G = (N, A)$, where N is a finite set of nodes representing spatial locations such as rooms or corridor segments, and $A \subseteq N \times N$ is a set of directed arcs representing walkable connections such as doors or passages. The graph is assumed to be simple, meaning that between any two adjacent nodes there exists at most one directed arc in each direction. If multiple distinct physical connections exist between the same pair of locations, they are modeled by introducing intermediate (fictitious) nodes, thereby preserving the simplicity of the graph.

Each directed arc $(u, v) \in A$ is associated with a nonnegative geometric length $d(u, v) \geq 0$ and a traversal time $w(u, v) \geq 0$. The traversal time $w(u, v)$ is derived from the average evacuation speed $\bar{v}(u, v)$ along arc (u, v), its geometric length $d(u, v)$, and additional infrastructural factors affecting movement (e.g., stairways or other impediments). In particular, it may be expressed as $d(u, v)/\bar{v}(u, v)$, augmented by additive penalty terms when such features are present.

Evacuation decisions are made over a discrete time horizon $T = \{0, 1, \dots, T_{max}\}$. Let $O_t \subseteq N$ denote the set of origin nodes at time $t \in T$, i.e., the nodes currently occupied by individuals who still require evacuation. Let $D \subseteq N$ denote the set of designated safe exit nodes. We assume that individuals requiring evacuation are never located at exit nodes so that $O_t \cap D = \emptyset$ for all $t \in T$.

For each origin node $s \in O_t$, let $\mathcal{P}_s$ denote the set of the first k shortest simple paths from s to any exit node $d \in D$, computed with respect to the nominal arc travel times $w(u, v)$. Each path $p \in \mathcal{P}_s$ is a simple (i.e., loopless) path, defined as an ordered sequence of adjacent distinct nodes $p = (v_1, v_2, \dots, v_l)$ with $v_1 = s$, $v_l \in D$ and such that each consecutive pair of nodes is connected by a directed arc, i.e., $(v_i, v_{i+1}) \in A$, for all $i = 1, \dots, k-1$. The nominal evacuation time τ^p of a path p is given by the additive arc-weight sum:

$$\tau^p = \sum_{i=1}^{l-1} w(v_i, v_{i+1}). \tag{1}$$

Since the globally shortest simple path is contained in $\mathcal{P}_s$, the nominal evacuation time achievable from origin s is given by: $\tau_s^* = \min_{p \in \mathcal{P}_s} \tau^p$.

To account for near-optimal yet practically relevant alternatives, we define set $\overline{\mathcal{P}}_s(\gamma)$ as the subset of γ-efficient evacuation paths within the enumerated set $\mathcal{P}_s$:

$$\overline{\mathcal{P}}_s(\gamma) = \{p \in \mathcal{P}_s \mid \tau^p \leq (1 + \gamma)\tau_s^*\}, \tag{2}$$

where $\gamma \geq 0$ is an efficiency tolerance parameter controlling the allowable deviation from the nominal evacuation time.

Let $\tau_{\max}^t$ be an externally imposed maximum admissible evacuation time dependent on the actual time period of evacuation $t \in T$. Then, we additionally require that

$$(1 + \gamma)\, \tau_s^* \leq \tau_{\max}^t, \ \forall t \in T. \tag{3}$$

The set $\overline{\mathcal{P}}_s(\gamma)$ can be approximated in practice using k-shortest simple-path algorithms under the nominal arc costs (see, e.g., [36]). This set defines the admissible routing alternatives considered by the evacuation strategies in this work.

3.2. Risk perception and unsafe nodes

We define *risk information* as information derived from an evolving *hazard process*, which captures the underlying harmful phenomenon (e.g., flood, fire, or smoke propagation) that affects the navigability of the environment over time. Moreover, we

define *situational awareness* as the spatial extent (local/reactive vs. global/anticipatory) over which agents perceive and incorporate hazard-related risk information during the discrete time horizon T. This concept captures the availability and use of dynamic risk data by an evacuee agent and is *not* intended to represent prior map familiarity or prior knowledge of the environment.

Let $\theta_{\text{unsafe}} \geq 0$ denote an unsafe-node threshold used in routing to exclude hazardous locations from evacuation paths. Risk is represented through a time-dependent node-level function $R_t(n) \in [0,1]$, where $n \in N$ and $t \in T$ denotes discrete evacuation time steps.

A node is considered unsafe at time $t \in T$ if

$$R_t(n) \geq \theta_{\text{unsafe}}. \tag{4}$$

Accordingly, we define the time-dependent set of unsafe nodes $B_t \subseteq N$ as

$$B_t = \{ n \in N \mid R_t(n) \geq \theta_{\text{unsafe}} \}, \tag{5}$$

where $t \in T$. Unsafe nodes are temporarily treated as blocked, i.e., forbidden for traversal. A γ-efficient path $p \in \overline{\mathcal{P}}_s(\gamma)$ is considered safe at time $t \in T$ if and only if all its nodes belong to the feasible induced subgraph $G_t = G[N \setminus B_t]$.

Finally, *agile routing* denotes a strategy in which an agent evaluates candidate evacuation paths based on their evacuation betweenness centrality within the set of γ-efficient paths $\overline{\mathcal{P}}_s(\gamma)$ (Eq. (2)). At each time $t \in T$, the agent selects a path from the feasible subset $\overline{\mathcal{P}}_s(\gamma; G_t)$ by excluding those that traverse unsafe nodes; if this set is empty, the selection is performed over the full set $\overline{\mathcal{P}}_s(\gamma)$ according to the same centrality-based criterion.

Strategy 1: Efficient routing (baseline)

Let $\overline{\mathcal{P}}_s(\gamma; G_t)$ denote the subset of γ-efficient evacuation paths from each origin node $s \in O_t$ that remain feasible in the time-dependent induced subgraph G_t, i.e., $\overline{\mathcal{P}}_s(\gamma; G_t) = \{ p \in \overline{\mathcal{P}}_s(\gamma) \mid V(p) \cap B_t = \emptyset \}$, where, for a path $p = (v_1, \ldots, v_l)$, $V(p) := \{v_1, \ldots, v_l\}$ denotes the set of vertices visited along p.

Under the efficient routing baseline, at each time $t \in T$, an agent located at node s chooses a G_t-feasible γ-efficient path minimizing the nominal evacuation time:

$$p_s^{\text{eff}}(t) = \arg \min_{p \in \overline{\mathcal{P}}_s(\gamma, G_t)} \tau^p. \tag{6}$$

Path $p_s^{\text{eff}}(t)$ is the minimum–travel-time path within the feasible set of γ-efficient evacuation paths in G_t at time $t \in T$. More precisely, it corresponds to the nominally shortest path among the admissible paths in $\overline{\mathcal{P}}_s(\gamma; G_t)$.

If $\overline{\mathcal{P}}_s(\gamma; G_t) = \emptyset$, no safe γ-efficient path is available at time t under the current blocking constraint. In that case, to preserve progress toward an exit when the subgraph G_t becomes disconnected, the agent applies the fallback rule of selecting the best path from $\overline{\mathcal{P}}_s(\gamma)$:

$$p_s^{\text{eff,rel}}(t) = \arg \min_{p \in \overline{\mathcal{P}}_s(\gamma)} \tau^p. \tag{7}$$

Strategy 2: Agile (evacuation betweenness centrality–aware) routing

The agile path selection strategy prioritizes structurally adaptable (agile) routes within the set of feasible γ-efficient evacuation paths $\overline{\mathcal{P}}_s(\gamma; G_t)$.

The evacuation betweenness centrality measure of a node $n \in N$ at time $t \in T$ is defined as the average frequency with which n appears as an intermediate vertex (i.e., excluding end vertices) along the feasible γ-efficient paths in $\overline{\mathcal{P}}_s(\gamma; G_t)$, aggregated over all currently active origin nodes O_t, similarly to [14]:

$$C_{EB,t}(n) = \sum_{\substack{s \in O_t \setminus \{n\} \\ \overline{\mathcal{P}}_s(\gamma, G_t) \neq \emptyset}} \frac{1}{\left| \overline{\mathcal{P}}_s(\gamma, G_t) \right|} \sum_{p \in \overline{\mathcal{P}}_s(\gamma, G_t)} \mathbf{1}_{\{ n \in V(p) \setminus \{s, d(p)\} \}}, \quad \forall n \in N \setminus D, \tag{8}$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function.

Each candidate path $p \in \overline{\mathcal{P}}_s(\gamma, G_t)$ is then assigned an (evacuation betweenness) agility value $\Delta_{EB,t}(p)$ by combining the evacuation betweenness centralities of its intermediate nodes:

$$\Delta_{EB,t}(p) = \prod_{n \in V(p) \setminus \{s, d(p)\}} C_{EB,t}(n). \tag{9}$$

Accordingly, at time $t \in T$, the agile path selection strategy chooses a G_t-feasible γ-efficient path that maximizes the evacuation betweenness agility value $\Delta_{EB,t}$:

$$p_s^{\text{ag}}(t) = \arg \max_{p \in \overline{\mathcal{P}}_s(\gamma, G_t)} \Delta_{EB,t}(p). \tag{10}$$

If $\overline{\mathcal{P}}_s(\gamma; G_t) = \emptyset$, the agile path selection strategy selects the path with the highest agility value from the γ-efficient path set $\overline{\mathcal{P}}_s(\gamma)$:

$$p_s^{\text{ag,rel}}(t) = \arg \max_{p \in \overline{\mathcal{P}}_s(\gamma)} \Delta_{EB,t}(p). \tag{11}$$

Table 1
Summary of routing-related terminology.

Term	Definition
Efficient routing	Path selection strategy that minimizes nominal evacuation time within the feasible γ-efficient path set (Eq. (6)).
Agile routing	Path selection strategy that prioritizes structurally adaptable routes by maximizing evacuation betweenness centrality within the γ-efficient path set (Eq. (10)).
γ-efficient paths	Set of near-optimal paths whose travel time does not exceed $(1+\gamma)\tau^*$ (Eq. (2)).
Situational awareness	Spatial extent over which agents perceive and incorporate risk information (low: local/reactive; high: global/anticipatory).
Fallback mechanism	Routing rule applied when no feasible path exists, selecting the best available path without safety filtering (Eqs. (7) and (11)).
Unsafe nodes	Nodes whose risk exceeds the threshold θ_{unsafe}.

This fallback preserves the agility-oriented ranking even when no currently safe γ-efficient path remains available.

The proposed agile path selection strategy is a simplified adaptation of the approach proposed in [14]. While the original formulation includes coordination based on evacuation centrality and dissimilar-path selection mechanisms, the present implementation focuses on the evacuation betweenness centrality (Eq. (8)) and the corresponding path agility (Eq. (9)).

3.3. Awareness-specific rerouting rules

In the proposed routing approach, awareness determines the spatial extent to which hazard-related risk information is perceived and used for rerouting decisions. Accordingly, we next introduce the low- and high-awareness regimes. Both regimes rely on the same unsafe-node definition (Eq. (4)) and on the γ-efficient evacuation paths set $\overline{\mathcal{P}}_s(\gamma)$ (Eq. (2)), but they differ in how much of the currently planned route is inspected for possible hazard-induced infeasibility.

Low-awareness (local and reactive behavior). Low-awareness agents behave reactively and monitor only the next node along the currently selected route. Let n_{next} denote the successor node to be visited along the selected path $p \in \overline{\mathcal{P}}_s(\gamma)$.

An agent evaluates the rerouting at time t whenever $R_t(n_{\text{next}}) \geq \theta_{\text{unsafe}}$. The agent then updates the blocked set B_t and recomputes the set $\overline{\mathcal{P}}_s(\gamma; G_t)$.

If $\overline{\mathcal{P}}_s(\gamma; G_t) \neq \varnothing$, the agent reroutes by applying the active routing strategy over this G_t-feasible subset. That is, the efficient path selection strategy selects the path with minimal nominal evacuation time, whereas the agile path selection strategy selects the path with the maximal agility value. Otherwise, if $\overline{\mathcal{P}}_s(\gamma; G_t) = \varnothing$, the agent acts on the fallback rule defined for the active routing strategy.

High-awareness (global and anticipatory behavior). High-awareness agents monitor the entire remaining planned route anticipatorily. A rerouting event is triggered whenever at least one future node along the current path becomes unsafe, i.e., whenever the currently planned path p satisfies $\{v_2, \ldots, v_l\} \cap B_t \neq \emptyset$.

As long as $\overline{\mathcal{P}}_s(\gamma; G_t) \neq \varnothing$, the agent applies the active routing strategy over this admissible set. Thus, awareness affects the anticipation of hazard-induced infeasibility, but it does not replace the primary decision rule associated with the selected routing strategy. On the contrary, if $\overline{\mathcal{P}}_s(\gamma; G_t) = \varnothing$, the blocking constraint is temporarily relaxed and the agent follows the fallback rule of the active routing strategy.

For simplicity, Table 1 summarizes the main routing-related terms used in this work and their relationships.

3.4. Performance metrics

A wide range of performance indicators can be defined to evaluate evacuation strategies in hazardous and dynamic environments, covering efficiency, safety, and decision robustness. This section reports a compact set of *key performance indicators* (KPIs) consistent with the routing framework introduced above: (i) *evacuation efficiency*, captured by evacuation time metrics; and (ii) *risk-related characteristics* of the selected routes under different awareness and routing strategies. Risk-based indicators are reported exclusively for evaluation and analysis and do not define an additional optimization objective.

As discussed in Section 4.5, pedestrians are organized into groups (typically reflecting shared initial positions on the same node). Let $\mathcal{G}$ denote an evacuation group of agents, $|\mathcal{G}|$ its size, and T_i the evacuation completion time of agent $i \in \mathcal{G}$. At each time period $t \in T$, risk statistics are computed over the remaining evacuation path for each evacuation group $\mathcal{G}^{(t)}$, excluding their current node and the destination node.

Evacuation time metrics (primary objective). The primary indicator is the *maximum evacuation time* $\tau^*_{\max}$, defined as the time at which the last agent in the group reaches a safe exit node:

$$\tau^*_{\max}(\mathcal{G}) = \max_{i\in\mathcal{G}} T_i, \tag{12}$$

where T_i denotes the evacuation completion time of agent $i \in \mathcal{G}$. This metric highlights worst-case performance and bottlenecks affecting the slowest evacuees. Complementarily, the *average evacuation time* captures typical group behavior:

$$\bar{T}(\mathcal{G}) = \frac{1}{|\mathcal{G}|}\sum_{i\in\mathcal{G}} T_i. \tag{13}$$

In addition, robust distributional statistics are reported: the *minimum evacuation time* $T_{\min}$, the *median evacuation time* T_{50}, and the *90th percentile evacuation time* T_{90}:

$$T_{\min}(\mathcal{G}) = \min_{i\in\mathcal{G}} T_i, \tag{14}$$

$$T_{50}(\mathcal{G}) = \text{median}\{T_i \mid i \in \mathcal{G}\}, \tag{15}$$

$$T_{90}(\mathcal{G}) = \text{percentile}_{90}\{T_i \mid i \in \mathcal{G}\}. \tag{16}$$

Average remaining path cost (route efficiency proxy). To assess structural route efficiency independently of transient effects, we compute the *average remaining path cost* in terms of arc travel times. Let $p^{(t)}_{\mathcal{G}} = (v_1, \dots, v_{k_t})$ denote the remaining path assigned to group $\mathcal{G}$ at decision time $t \in T$, where v_1 is the current node and $v_{k_t} = d\left(p^{(t)}_{\mathcal{G}}\right)$ is the destination exit node. The group's remaining nominal path travel time $\tau\left(p^{(t)}_{\mathcal{G}}\right)$ is defined as:

$$\tau\left(p^{(t)}_{\mathcal{G}}\right) = \sum_{j=1}^{k_t-1} w(v_j, v_{j+1}), \tag{17}$$

where $w(\cdot,\cdot)$ is the traversal-time cost used by the routing layer, as defined in Section 3.1, and k_t denotes the number of nodes in the remaining path $p^{(t)}_{\mathcal{G}}$ at time $t \in T$.

The *average remaining path travel time* $\bar{\tau}(\mathcal{G})$ is computed across all decision times in the evacuation horizon $\mathcal{T}_{\mathcal{G}}$ of group $\mathcal{G}$ as:

$$\bar{\tau}(\mathcal{G}) = \frac{1}{|\mathcal{T}_{\mathcal{G}}|}\sum_{t\in\mathcal{T}_{\mathcal{G}}} \tau\left(p^{(t)}_{\mathcal{G}}\right), \tag{18}$$

where $\tau\left(p^{(t)}_{\mathcal{G}}\right)$ denotes the remaining nominal path travel time at decision time t. This metric supports direct comparison between strategies that favor shorter nominal routes and those that prioritize structurally flexible alternatives, especially under frequent rerouting.

Remaining-path risk statistics (route safety characterization). Risk statistics are computed over the intermediate nodes of the path, i.e., excluding the current node v_1 and the destination node v_{k_t}.

The *mean remaining-path risk* (Mean RPR) is defined as:

$$\mu_R\left(p^{(t)}_{\mathcal{G}}\right) = \frac{1}{k_t - 2}\sum_{j=2}^{k_t-1} R_t(v_j), \qquad \text{for } 3 \le k_t \le |N|. \tag{19}$$

The corresponding *remaining-path risk variance* (RPR Var.) is defined as

$$\text{Var}_R\left(p^{(t)}_{\mathcal{G}}\right) = \frac{1}{k_t - 2}\sum_{j=2}^{k_t-1} \left(R_t(v_j) - \mu_R\left(p^{(t)}_{\mathcal{G}}\right)\right)^2, \qquad \text{for } 3 \le k_t \le |N|. \tag{20}$$

For $k_t = 2$, this quantity is defined as zero by convention, since the remaining path contains no intermediate nodes.

Low variance indicates that risk is relatively uniform along the remaining route, while high variance indicates localized high-risk segments that may trigger rerouting even if the mean risk is moderate.

For interpretability, we also report the *maximum remaining-path risk*:

$$R_{\max}\left(p^{(t)}_{\mathcal{G}}\right) = \max_{j\in\{2,\dots,k_t-1\}} R_t(v_j), \qquad \text{for } 3 \le k_t \le |N|. \tag{21}$$

As in the previous case, for $k_t = 2$, this quantity is defined as zero by convention, since the remaining path contains no intermediate nodes. This metric directly relates to threshold-based decisions (e.g., whether any upcoming node satisfies Eq. (4)).

Rerouting and decision robustness (optional behavioral KPIs). To quantify how reactive or stable each strategy is, the simulation simultaneously records: (i) the *number of rerouting events* for group $\mathcal{G}$, denoted $N_{\text{re}}(\mathcal{G})$, counting how many times the planned path is updated; and (ii) the *path-change rate*, defined as $N_{\text{re}}(\mathcal{G})/|\mathcal{T}_{\mathcal{G}}|$. These indicators help explain performance differences when hazards evolve quickly and strategies frequently revise decisions.

Aggregation across groups and scenarios. Although KPIs are computed at the group level, results are reported aggregated by routing strategy and situational awareness level. For each configuration, the presented values correspond to the average of the group-level metrics across all groups operating under identical conditions. Dispersion measures (e.g., standard deviation) can be added to reflect stochastic variability across simulation runs.

4. Simulation architecture

This study relies on a modular simulation framework designed to evaluate evacuation performance under dynamic environmental hazards and heterogeneous levels of situational awareness. The framework consists of three main modules that operate on a shared graph-based environment: (i) a *Risk Simulation Module* that models the temporal evolution of hazards, (ii) a *Routing Module* responsible for computing evacuation paths based on agent awareness and predefined routing strategies, and (iii) an *Agent Simulation Module* that executes the pedestrian movement dynamics.

This separation of concerns ensures clarity, extensibility, and the ability to modify any of the three processes without affecting the others. Such modular architectures have become a standard in recent evacuation simulation frameworks, allowing for flexible integration of physical and cognitive models [15,16].

4.1. Graph-based environment representation

As introduced in Section 3.1, the simulation environment is formalized as a directed graph $G = (N, A)$. Nodes N represent spatial entities such as rooms, corridor segments, or safe exits, while arcs A denote walkable connections between them. Each arc $(u, v) \in A$ is associated with a geometric length $d(u, v) \geq 0$ (e.g., Euclidean distance in the embedded environment). For routing and motion, this geometric information is mapped to traversal-time costs via the nominal arc-cost model introduced in Section 3.1. In particular, each arc $(u, v) \in A$ is assigned a traversal-time cost $w(u, v)$, which depends on its geometric length $d(u, v)$ and on infrastructural properties affecting movement (e.g., stair penalties and assumed walking speeds).

This graph-based representation constitutes a shared spatial abstraction across the entire framework. It is directly exploited by the Routing Module for path selection, by the Risk Simulation Module for hazard propagation, and by the Agent Simulation Module for movement and coordination, ensuring a unified interpretation of geometry, connectivity, and neighborhood relations. Graph-based modeling of indoor environments has been widely adopted for its computational efficiency and for enabling analytical approaches based on network centrality [13,14].

At the start of each simulation, the environment provides only static information to the remaining modules: the graph G topology, the geometric embedding (i.e., node coordinates and/or arc lengths), the initial spatial positions of all agents, and the initial configuration of hazards (i.e., ignition nodes and their associated risk values). No risk propagation or routing decisions are performed at this stage. Instead, the environment acts as a structural substrate upon which dynamic processes — risk evolution, routing decisions, and pedestrian motion — are subsequently executed.

4.2. Risk modeling

The *Risk Simulation Module* is responsible for computing the temporal and spatial evolution of hazards prior to, or in parallel with, the agent-based simulation. It operates directly on the navigation graph G and receives as input: (i) the initial risk configuration, (ii) model parameters such as the activation threshold θ_{act} and the propagation probability, and (iii) the number of frames to simulate. Its output is a complete time-indexed *risk evolution map* $R_t(v)$, where $v \in N$ denotes a node in the graph, t a discrete time step, and $R_t(v) \in [0, 1]$ the risk value associated with node v at time t.

An activation threshold $\theta_{\text{act}} \geq 0$ is used internally by the hazard model to activate new hazard sources during the evacuation simulation. The time-dependent set of unsafe nodes is part of the Routing Module itself and is applied uniformly across all routing strategies and situational awareness levels. Awareness influences when and over which spatial horizon rerouting decisions are triggered, while the unsafe-node feasibility logic remains identical across configurations. When no feasible γ-efficient path remains available under the current blocking constraint, the fallback behavior consists of temporarily relaxing the feasibility filter and applying the relaxed decision rule associated with the active routing strategy.

The propagation mechanism is based on a distance-attenuated diffusion process. When a node exceeds the activation threshold θ_{act}, it becomes an *active emitter* and transfers a fraction of its risk to neighboring nodes. For the experimental evaluation, we adopt a simple, structured propagation scheme in which first-order neighbors receive one third of the emitter's current risk, while second-order neighbors receive one ninth. These coefficients enforce a monotonic decay with distance in the graph, providing a controlled and interpretable approximation of diminishing hazard intensity.

This propagation component is implemented as a replaceable module. Consequently, the proposed architecture supports alternative formulations — such as exponential decay, nonlinear diffusion, anisotropic spreading, or empirically calibrated hazard dynamics — without modifications to the routing, decision, and evaluation layers. This design choice facilitates integration with perception-based risk models in which hazard propagation is coupled to visibility and awareness dynamics [2,3].

In the current implementation, hazard dynamics are modeled as a time-indexed risk field on the navigation graph G and are computed independently of pedestrian motion. This abstraction enables controlled comparisons across routing and awareness configurations; future work will couple hazard evolution to environment-state variables (e.g., ventilation/smoke transport, door states, and visibility) and, where relevant, to agent–hazard interactions.

The result of the risk simulation is a time-indexed risk map $R_t(\cdot)$ that is later used by both the Routing Module and the Agent Simulation Module. As a consequence, risk levels may increase over time according to the hazard model, yielding progressively constrained evacuation scenarios. For clarity, the routing layer classifies nodes as unsafe by comparing $R_t(n)$, for all $n \in N$ against the unsafe-node threshold θ_{unsafe} defined in Section 3.2, which is conceptually distinct from the activation threshold θ_{act} used by the propagation model.

4.3. Objective prioritization and parameter setting

Primary optimization objective. When evaluating evacuation strategies, a fundamental trade-off arises between maintaining short evacuation routes and preserving adaptability under dynamic hazards. In the proposed routing approach, the admissible solution space is first defined by nominal evacuation time τ_s^* and the associated set of γ-efficient evacuation paths $\overline{P}_s(\gamma; G_t)$.

Hazard-related risk is not treated as a competing optimization objective. Instead, it acts as a feasibility filter and rerouting trigger: nodes satisfying $R_t(n) \geq \theta_{\text{unsafe}}$ are considered unsafe (Eq. (4)) and are excluded through the blocked set B_t (Eq. (5)), thereby restricting the admissible set of candidate evacuation paths.

Within the resulting feasible subset, the primary decision rule depends on the active routing strategy. Under the efficient strategy, path selection minimizes nominal evacuation time (Eq. (6)); under the agile strategy, path selection maximizes the dynamic centrality-based score (Eq. (10)). Remaining-path risk statistics are retained for evaluation and interpretation, rather than as a general optimization criterion.

If an explicit tie-breaking rule is required among paths that are equivalent under the active routing criterion, remaining-path risk may be used only as a tertiary preference indicator. In that case, the preferred path is the one with lower mean remaining-path risk, with lower risk variance used only as a subsequent tie-breaker.

Trade-off between efficiency and centrality. The role of centrality is explicitly controlled by the tolerance parameter γ. Even under the agile strategy, paths are considered for centrality-based ranking only if they satisfy the near-optimality condition in Eq. (2). Consequently, a path with high structural centrality but substantially larger nominal evacuation time than the shortest available route is never considered admissible.

Centrality therefore does not replace the travel-time constraint; rather, it provides a secondary ranking principle within the γ-efficient path set. In contrast, risk information does not define an additional optimization dimension, but determines whether a candidate path remains feasible and whether rerouting must be triggered under the corresponding awareness regime.

This separation preserves comparability between routing strategies: nominal evacuation time defines the γ-efficient evacuation path set, centrality differentiates alternative agile paths within that set, and risk governs feasibility monitoring and rerouting activation.

Parameter configuration and justification. All routing-relevant parameters are specified per scenario via external configuration files. In the present implementation, the key parameters governing routing behavior are: (i) the unsafe-node threshold θ_{unsafe}, which determines node feasibility (Eq. (4)), and (ii) the tolerance parameter γ, which controls admissible deviations from the minimum nominal evacuation time (Eq. (2)). The activation threshold θ_{act} is used only within the hazard propagation model (Section 4.2) and does not affect routing decisions directly.

The values reported in this paper were selected through a pilot sweep of the parameters θ_{unsafe} and γ (parameter ranges are reported in Appendix A.1). The retained configurations (i) avoided degenerate behavior (e.g., constant rerouting or frequent fallbacks due to extensive node blocking), (ii) preserved a clear separation between routing strategies in evacuation-time and route-characterization metrics, and (iii) did not induce unbounded detours beyond the $(1+\gamma)$ travel time constraint (Eq. (2)). The selected values are not claimed to be optimal nor calibrated to a specific facility; rather, they provide stable and interpretable comparisons under the considered hazard dynamics.

Importantly, no parameter is hard-coded in the model logic: all routing-relevant parameters are defined externally through scenario-specific configuration files. For reproducibility and transparency, the complete set of parameter values used in each experiment — including routing strategy, awareness level, θ_{unsafe}, γ, and risk dynamics — is documented in the configuration files listed in Appendix A.1.

The current formulation adopts static values for γ and for the maximum admissible evacuation time $\tau_{\max}^t$ (Eq. (3)) to ensure controlled and interpretable comparisons across routing strategies and awareness regimes. However, in rapidly evolving hazard scenarios, this static parameterization may limit the ability to anticipate risk propagation. In particular, a fixed γ may restrict the diversity of admissible routes, potentially delaying proactive rerouting, while a static $\tau_{\max}^t$ may not adequately reflect the shrinking safe time horizon induced by increasing risk levels. A joint dynamic adaptation of these parameters would allow the routing process to progressively shift from an efficiency-oriented regime toward a safety- and resilience-oriented regime as environmental conditions deteriorate. While such adaptive mechanisms are beyond the scope of the present study, they represent a promising direction for future work, particularly in scenarios characterized by fast hazard propagation or limited route redundancy.

4.4. Routing module

The *Routing Module* acts as the decision-making component responsible for providing evacuation paths to agent groups during the simulation. Rather than executing movement or simulating pedestrian dynamics, its role is limited to computing and updating routes when requested by the Agent Simulation Module.

Routing decisions are based on three inputs: (i) the shared graph topology, (ii) the time-dependent risk map $\mathcal{R}_t(v)$ generated by the Risk Simulation Module (Section 4.2), and (iii) group-specific attributes such as the assigned routing strategy and situational awareness level, as defined in Sections Section 3. This design reflects recent evidence that routing and awareness interact dynamically: awareness affects how agents interpret and react to environmental information, while routing logic determines the structural feasibility of those reactions [9,10].

4.5. Agent simulation

The *Agent Simulation Module* executes the pedestrian dynamics. At initialization time, it receives the environment topology, the risk evolution map generated by the Risk Simulation Module, and the initial spatial positions of all agents.

Once the simulation begins, agents start their movement along their assigned routes using JuPedSim's microscopic motion model. Decision making processes during the pre-evacuation phase (i.e. the behavior of occupants of the environment before the purposive movement towards an exit) is actually quite complex (see, e.g., [37]), but for the sake of evaluating the route guidance mechanism the experiments focused on the evacuation phase assuming a simplistic simultaneous start of the evacuation phase for all pedestrians. In all simulation experiments, pedestrian movement was modeled using JuPedSim's *CollisionFreeSpeedModel*, which provides realistic microscopic dynamics while ensuring stable interactions among agents and obstacles. However, since the proposed framework is fully modular, the same simulation setup can be executed using any of JuPedSim's alternative motion models without modification to the routing or risk modules. This aligns with recent hybrid frameworks that integrate physics-based motion with cognitive or awareness layers to capture adaptive pedestrian behavior [15,16].

Based on the initial positions, agents are automatically grouped: all agents originating from the same starting node are assigned to the same group and, as a consequence, share the same routing strategy and level of situational awareness. This grouping mechanism reflects the empirical tendency of evacuees in close spatial proximity to exhibit coordinated behavior [2]. Route choice and rerouting decisions are performed at the group level through a leader-based coordination mechanism. Before describing the adopted mechanism, let us first consider that groups generally represent the preponderant part of the set of pedestrians in a given situation: one of the first works focusing on the impact of groups in pedestrian dynamics reported a proportion of pedestrians belonging to groups between 55 and 70% of the considered population [38], with following works reporting even higher values for the group share (see, e.g., [39]). Understandably, modelers have started considering the influence of groups on pedestrian dynamics, but these efforts proposed modifications to operational level mechanisms of the simulation model [40]. The approach adopted for the present work defines instead mechanisms related to tactical level decisions. At any given time, a *group leader* is defined as the agent that is furthest along the group's planned evacuation path. This agent is considered to have the most relevant spatial perspective and is therefore responsible for monitoring hazard-related risk along the currently assigned route according to the group's situational-awareness regime. Under low-awareness conditions, the leader evaluates only the next node to be visited and triggers rerouting when that node exceeds the unsafe threshold. Under high-awareness conditions, the leader inspects the entire remaining planned route and requests rerouting whenever any future node becomes unsafe. In both cases, the leader delegates the computation of the new route to the routing module, which applies the active routing strategy over the currently feasible set of candidate paths.

Leadership is dynamic, and the change of leader is related to rerouting decisions: as agents progress, congestion patterns shift, or speed differences emerge, the agent furthest along the path may change. When this occurs, leadership automatically transfers to the new front-most agent, ensuring that decisions consistently rely on the agent with the most advanced local knowledge. This leader-based adaptation mechanism parallels findings in recent behavioral studies showing that local coordination and group imitation play a critical role in maintaining efficient flow under stress [10].

If the routing module returns a safer alternative path, all members of the group update their trajectories simultaneously, preserving coordinated movement. If no preferable alternative exists, the group continues on its current route. This leader-driven decision-making approach enhances behavioral realism and improves computational efficiency by avoiding redundant per-agent route evaluations.

Through the interplay of risk evolution, routing decisions, leadership dynamics, and microscopic movement, the Agent Simulation Module captures adaptive and coherent evacuation behavior under dynamically changing environmental conditions.

4.6. Framework overview and decision-making levels

Fig. 1 summarizes the modular structure of the proposed simulation architecture and clarifies the interaction between its components. The figure highlights how the graph-based environment representation serves as the common substrate shared by all modules, while risk evolution, routing decisions, and pedestrian dynamics are handled by distinct but interconnected processes.

From a conceptual perspective, the architecture can be interpreted through the classical three-level decision-making framework commonly adopted in pedestrian dynamics and evacuation modeling [27]. The *strategic level* is represented by the initial configuration of the graph-based environment, including the spatial layout, initial agent locations, and initial hazard conditions. These elements define the evacuation context and remain fixed throughout the simulation.

The *tactical level* corresponds to the Routing Module, where route selection and adaptation decisions are made based on routing strategies and situational awareness. Tactical decisions are triggered during the simulation when rerouting conditions arise and are informed by the time-dependent risk map produced by the auxiliary Risk Simulation Module.

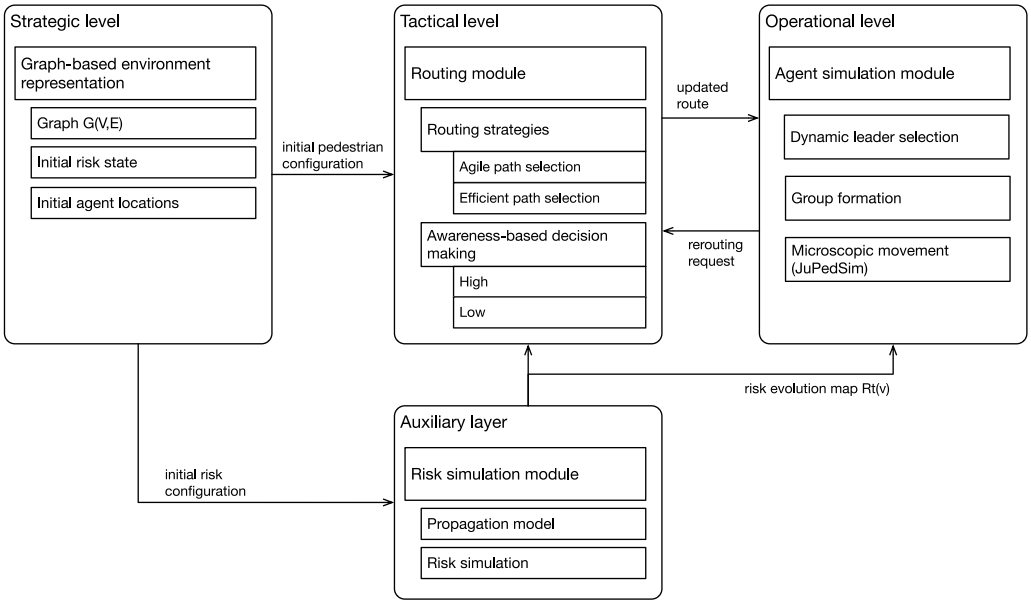

Fig. 1. Modular architecture of the evacuation simulation framework.

Finally, the *operational level* is implemented by the Agent Simulation Module, which governs the continuous execution of pedestrian movement through microscopic dynamics and group-level coordination. Operational behavior follows the routes determined at the tactical level while responding locally to congestion and interactions with other agents. This layered interpretation aligns with current hybrid or multi-scale simulation paradigms that link cognitive awareness, tactical decision-making, and physical movement into unified evacuation models [9,15]. What is indicated as *auxiliary layer* is essentially an environmental abstraction, in line with [41], that encapsulates risk simulation and propagation mechanisms.

By organizing the framework along these three decision-making levels, the proposed architecture achieves a clear separation of concerns while preserving the necessary interactions between modules. This structure supports modularity and extensibility and provides a coherent foundation for analyzing how routing strategies and situational awareness influence evacuation performance under dynamically evolving hazards.

5. Simulation environments and scenario design

This section introduces the simulation environments used in the study, designed to elicit non-trivial routing decisions under uncertainty. To meaningfully assess the influence of routing strategies and situational awareness, the environments must present agents with structurally complex scenarios where multiple viable alternatives exist and where decision-making has tangible consequences. Overly simple layouts, such as single-corridor configurations with limited intersections, are deliberately excluded, as they funnel agents through dominant routes, trivialize routing choices, and obscure the behavioral differences induced by alternative decision mechanisms. These situations can hardly benefit from adaptive evacuation guidance systems, since pedestrians have very limited routing alternatives; in such layouts, capacity constraints at passages and bottlenecks dominate evacuation dynamics. The selected environments, by contrast, ensure that routing behavior remains a critical determinant of evacuation performance. They constitute the foundation for all experimental analyses discussed in the following section.

Fig. 2 illustrates an example of a minimal environment. In this configuration, agents have no real opportunity to demonstrate adaptive behavior: the optimal path is immediately apparent, rerouting options are negligible, and situational awareness provides little practical benefit. While such environments may be useful for algorithm validation or debugging, they are insufficient for evaluating decision-making under dynamic risk conditions.

In contrast, the environments selected for this work incorporate substantial topological and spatial complexity. They include branching corridors, loops, even potential dead ends; there are multiple decision points, creating conditions in which agents must continuously evaluate alternative routes, react to evolving hazard sources, and potentially reroute under pressure. Such features are essential to expose differences between routing strategies and to rigorously assess the impact of situational awareness on evacuation performance.

Figs. 3 and 4 present the two representative layouts used throughout the experimental evaluation: a theme park-inspired environment and a cruise ship deck layout. Although both environments exhibit comparable levels of structural richness, they represent distinct classes of evacuation scenarios.

The theme park environment (Fig. 3) is characterized by wide corridors, open intersections, and multiple converging routes, resulting in high navigational ambiguity and a strong potential for flow divergence and localized congestion. This layout allows

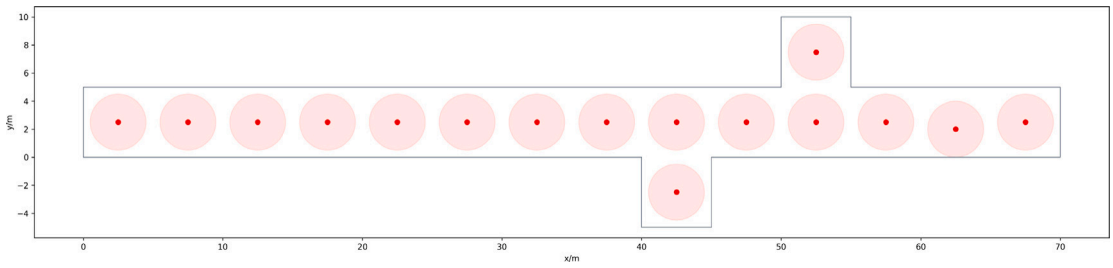

Fig. 2. Simplified linear environment. Agents have only one dominant route, limiting opportunities for decision-making or adaptation.

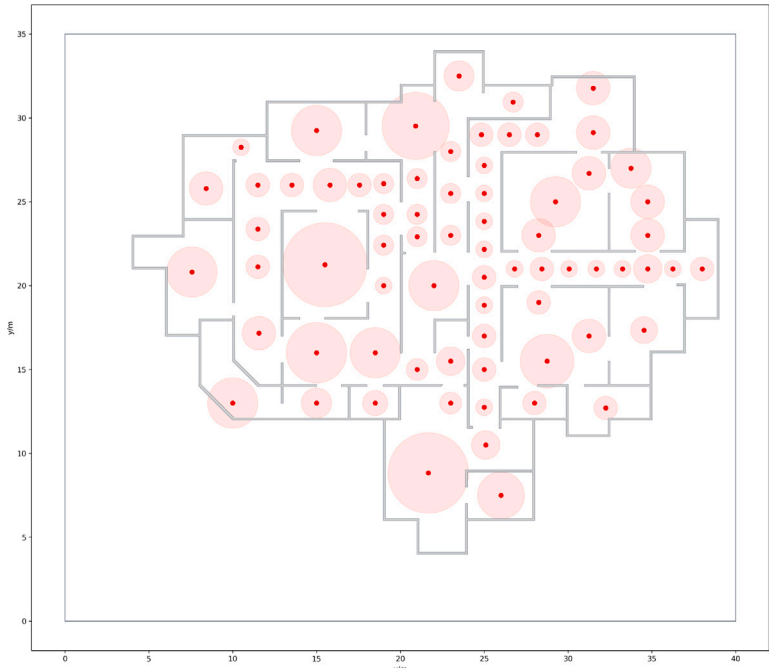

Fig. 3. Theme park simulation environment layout.

agents to choose among several viable paths, making it well suited to study how routing strategies distribute agents across space and how awareness influences route selection in open environments.

In contrast, the cruise ship layout (Fig. 4) imposes stronger spatial constraints through narrow corridors, compartmentalized areas, and pronounced structural bottlenecks. These characteristics amplify the effects of congestion and hazard propagation, forcing agents to navigate confined spaces where routing decisions have more immediate consequences. As a result, this environment is particularly suitable for analyzing the robustness of routing strategies and the benefits of timely hazard perception in constrained settings.

By incorporating these two complementary environments, this study ensures that routing decisions are not trivialized and that agents are required to balance efficiency and safety in structurally rich spaces. Together, these layouts enable robust testing of evacuation behavior under varying levels of uncertainty, pressure, and environmental constraint, supporting the validity and applicability of the proposed evacuation framework.

6. Experimental evaluation

This section presents the experimental evaluation of the proposed evacuation framework, focusing on the interaction between routing strategies and situational awareness under dynamic risk conditions. All simulations were conducted in the environments described in Section 5, which provided sufficient structural complexity to induce non-trivial decision-making and adaptive behavior.

The evaluation proceeded in three stages. First, a representative scenario involving dynamic hazard expansion was analyzed in detail to illustrate how awareness and routing logic affected rerouting behavior as risks evolved. Second, the same combinations of routing strategy and awareness level were systematically tested in the two complex layouts (theme park and cruise ship), supported

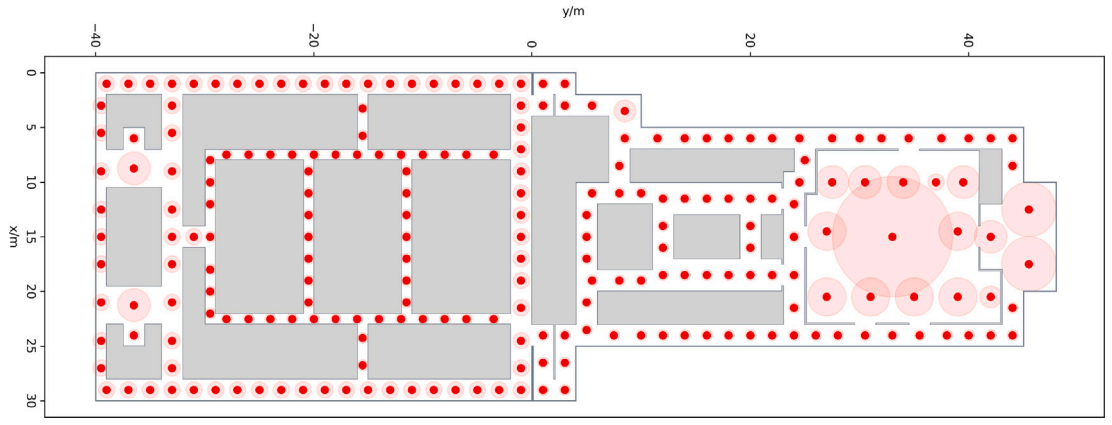

Fig. 4. Cruise ship simulation environment layout.

Table 2
Performance metrics for different routing strategies and awareness levels in the Cruise Ship scenario with dynamic hazard expansion. RPR denotes remaining-path risk; the remaining-path cost is computed as the sum of edge costs along the remaining route.

Group	Configuration	Max Evac. Time (s)	Avg Evac. Time (s)	Avg remaining-path cost	Mean RPR	RPR Var.
Blue Group	Centrality, High-Awareness	125.10	120.54	67.16	0.0012	0.0002
Blue Group	Centrality, Low-Awareness	129.09	126.31	61.22	0.0080	0.0023
Red Group	Efficient, High-Awareness	135.39	132.08	75.52	0.0004	0.0002
Red Group	Efficient, Low-Awareness	162.09	158.07	76.29	0.0326	0.0136

by heatmaps and aggregated performance metrics. Third, an auxiliary experiment explored the effects of increased agent density and exit availability on congestion patterns and route utilization. A comparative summary concludes the section, highlighting dominant trends across all scenarios.

6.1. Evacuation with dynamic hazard expansion

To illustrate how routing logic and awareness interact in a dynamic environment, we examined in detail a representative case involving two groups evacuating simultaneously. The environment began in a stable state, but risks evolved as the simulation progressed, forcing agents to react to changing threats.

Fig. 5 provides a visual overview of the initial conditions and subsequent risk evolution in this scenario. At the beginning of the simulation, two groups of five agents were located in the upper part of the environment, corresponding to two distinct starting locations. Two evacuation exits were highlighted in light pink near the lower corners of the layout and were identified by the letters A and B. In addition, several areas of the environment were overlaid with a color gradient ranging from light pink to dark purple, representing different levels of environmental risk. Lighter shades indicated low-risk regions, while darker purple areas corresponded to risk values exceeding the defined unsafe-node threshold θ_{unsafe}.

The figure also illustrates the two key phases of this scenario. During the initial stage (Fig. 5(a)), a localized hazard formed on the right side of the environment. Agents unaffected by this early danger proceeded without major deviation, while others began to alter their planned paths. Later in the simulation (Fig. 5(b)), a second, unexpected hazard emerged elsewhere in the environment, compelling additional rerouting depending on each agent's awareness level and routing strategy.

This case highlighted performance differences across routing strategies and awareness levels. Under low-awareness conditions, agents following efficient-path routing tended to persist longer on routes that later became unsafe, resulting in higher Mean RPR and larger variability. In contrast, higher awareness and structurally robust routing logic promoted earlier route switching and consistently lower remaining-path risk.

Table 2 reports the main outcome indicators for each group and configuration. In line with the rest of this section, risk was reported using remaining-path risk metrics: Mean RPR and its variability (RPR Var.). The *Avg remaining-path cost* corresponded to the sum of edge costs (i.e., traversal-time costs) along the remaining route at decision points. Full recordings are available.[3]

Overall, the results were consistent with the central hypothesis of this study: effective evacuation under uncertainty was best achieved by combining routing heuristics that avoid structurally fragile paths with proactive awareness mechanisms that enabled timely rerouting when hazards evolved.

[3] Dynamic hazard expansion recordings: Low-awareness (https://vimeo.com/1094342872); High-awareness (https://vimeo.com/1094342863).

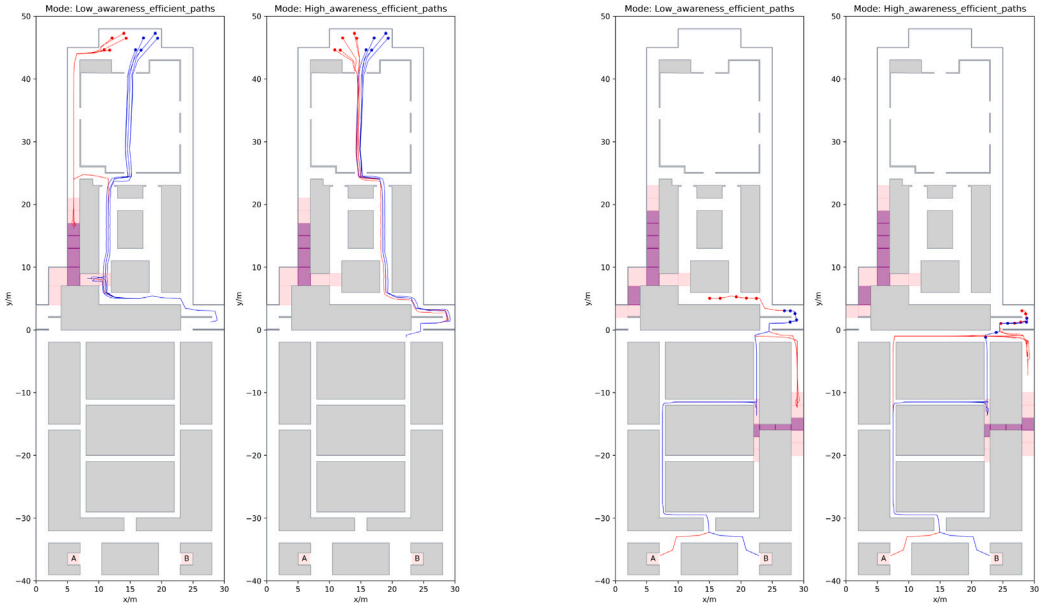

(a) Early stage: localized hazard begins to spread (purple).

(b) Late stage: new hazard emerges, changing risk landscape.

Fig. 5. Evolution of risk zones and agent behavior during evacuation.

It is worth noting that in some scenarios the efficient and centrality-aware strategies could converge to identical routing decisions. This typically occurred when the γ-efficient evacuation path set contained structurally similar paths, or when hazard evolution did not generate meaningful rerouting opportunities. In contrast, as illustrated in the present scenario with dynamic hazard expansion, differences between strategies emerged when late or unexpected hazards altered the risk landscape, causing alternative near-optimal paths to diverge in their structural centrality and rerouting potential.

6.2. Scenario-based evaluation across representative environments

To assess robustness across spatially distinct settings, we evaluated the same routing strategies and situational-awareness levels in the representative environments introduced in Section 5: Corridor, Theme Park, and Cruise Ship. For each layout, the spatial structure was kept fixed, while scenario variability was introduced through the initial agent distribution, group sizes, and the spatiotemporal emergence of dynamic hazards. Rather than designing stress tests under extreme congestion, these experiments examined how routing logic and situational awareness behaved in environments with different structural properties, where congestion and hazard interactions arose naturally.

For each environment, a set of scenarios was defined by varying initial conditions and hazard realizations while keeping the underlying layout unchanged. In total, ten scenarios were considered for the Corridor environment and eleven scenarios for each of the Theme Park and Cruise Ship environments. Each scenario was simulated once using a distinct random seed, thereby incorporating stochastic variability in pedestrian motion and interactions. Unless stated otherwise, tables report mean ± standard deviation aggregated over all scenarios for each environment, routing strategy, and awareness configuration.

Unless otherwise stated, all figures in this section correspond to a single representative scenario (hereafter referred to as the *representative case*) selected for the Theme Park and Cruise Ship environments. In contrast, all tables report aggregated metrics computed over the full set of scenarios. For the Cruise Ship environment, the example case introduced in Section 6.1 is additionally included, resulting in 12 evaluated cases (11 scenarios + 1 example case); all corresponding simulation runs are included in the reported aggregates. For the Corridor environment, we report only aggregated metrics, as the scenario-level visualization is less informative in this minimal layout.

To provide additional context on scenario variability, Table 3 reports summary statistics of the scenario configurations per environment. Specifically, we reported the mean ± standard deviation (across scenarios) of: the number of distinct initial nodes where agents were placed, the total number of agents, the number of evacuation target nodes, and the number of initial hazard nodes (risk sources). More extensive configuration information, including the full set of scenario parameters and configuration files, is provided in Appendix A.1.

Corridor environment. Table 4 reports aggregated metrics over all Corridor scenarios. As expected in this minimal layout, routing strategies yielded nearly identical evacuation-time and path-cost metrics, while differences across awareness levels primarily affected

Table 3
Scenario summary statistics (mean ± standard deviation) across all scenarios for each environment. "Initial hazard nodes" refer to the configured risk sources at the start of each scenario.

Environment	Avg N. initial agent nodes	Avg N. agents	Avg N. targets	Avg N. initial hazard nodes
Cruise Ship	3.25 ± 1.48	17.92 ± 9.50	2.00 ± 0.00	4.58 ± 1.66
Theme Park	3.27 ± 0.96	23.27 ± 6.63	1.55 ± 0.50	3.00 ± 2.04
Corridor	1.30 ± 0.64	26.50 ± 13.61	1.50 ± 0.50	1.40 ± 1.11

Table 4
Mean ± standard deviation of performance metrics aggregated over all Corridor scenarios and simulation runs. In this minimal layout, centrality-based and cost-based routing often coincide due to limited route diversity.

Configuration	Max Evac. Time (s)	Avg Evac. Time (s)	Avg remaining-path cost (s)	Mean RPR	RPR Var.
Centrality, High-Awareness	36.30 ± 10.89	31.30 ± 8.94	6.72 ± 2.44	0.0025 ± 0.0061	0.0009 ± 0.0028
Centrality, Low-Awareness	36.65 ± 11.28	31.86 ± 9.68	6.62 ± 2.42	0.0085 ± 0.0140	0.0034 ± 0.0055
Efficient, High-Awareness	36.30 ± 10.89	31.30 ± 8.94	6.72 ± 2.44	0.0025 ± 0.0061	0.0009 ± 0.0028
Efficient, Low-Awareness	36.65 ± 11.28	31.86 ± 9.68	6.62 ± 2.42	0.0085 ± 0.0140	0.0034 ± 0.0055

Table 5
Mean ± standard deviation of performance metrics aggregated over all Theme Park scenarios and simulation runs.

Configuration	Max Evac. Time (s)	Avg Evac. Time (s)	Avg remaining-path cost (s)	Mean RPR	RPR Var.
Centrality, High-Awareness	24.48 ± 7.83	20.28 ± 6.26	6.44 ± 3.05	0.0291 ± 0.0222	0.0036 ± 0.0029
Centrality, Low-Awareness	24.78 ± 7.74	20.64 ± 6.22	6.55 ± 2.96	0.0331 ± 0.0276	0.0054 ± 0.0060
Efficient, High-Awareness	24.39 ± 7.75	20.17 ± 6.19	6.32 ± 2.98	0.0300 ± 0.0210	0.0038 ± 0.0028
Efficient, Low-Awareness	27.38 ± 10.69	22.98 ± 9.04	6.74 ± 3.67	0.0442 ± 0.0369	0.0104 ± 0.0122

risk-based indicators (Mean RPR and RPR Var.). This pattern highlighted the role of awareness-specific risk inspection in shaping rerouting behavior under limited route diversity.

When extending the simulations to the multi-group scenarios considered in the Theme Park and Cruise Ship environments, the visualization strategy was adapted with respect to the representation used in the dynamic hazard expansion case described in Section 6.1. The visualization of environmental risk remained unchanged: shaded areas ranging from light pink to dark purple represented increasing risk levels, with darker tones indicating values exceeding the defined unsafe-node threshold θ_{unsafe}. However, the representation of agents and flows was modified to improve readability in denser scenarios.

In these figures, initial agent locations were annotated with white numerical labels indicating the number of agents initially present at each location. During the simulation, path traces no longer represented individual trajectories but instead encoded the cumulative number of agents that traversed each path segment. Evacuation targets were explicitly marked with the label *"target"* in white text, allowing exits to be clearly distinguished from intermediate nodes. This visualization scheme provided a compact and interpretable representation of agent distribution, movement intensity, and interaction with the evolving risk field when multiple groups are present simultaneously.

As described in Section 5, the Theme Park layout offered wide corridors, open intersections, and multiple alternative routes, resulting in high structural redundancy. These properties made the layout particularly sensitive to flow divergence and localized congestion effects. Complete recordings for the representative case are available.[4]

Fig. 6 illustrates the spatial distribution of agents and the initial hazard sources for the representative Theme Park scenario under all routing and awareness configurations. To complement these qualitative visualizations, Table 5 reports aggregated performance metrics in terms of mean ± standard deviation computed over all Theme Park scenarios and simulation runs.

The aggregated results indicated that higher situational awareness consistently led to lower mean remaining-path risk (Mean RPR) and reduced variability, reflecting safer and more stable evacuation behavior. Under high-awareness conditions, differences between routing strategies were small and remained within the reported standard deviations, suggesting that the high structural redundancy of the Theme Park environment largely mitigated the impact of the routing heuristic.

Under low-awareness conditions, however, agile routing achieved lower Mean RPR together with lower evacuation times than the efficient-path strategy, reinforcing that the reduction in risk could be obtained without a meaningful time penalty.

This contrast became even clearer in the more constrained Cruise Ship environment considered next.

The Cruise Ship environment introduced strong spatial constraints in the form of narrow corridors, compartmentalized spaces, and structural bottlenecks, which amplified the effects of congestion and hazard propagation. As in the Theme Park environment, the

[4] Theme Park simulations: Low-Awareness + Efficient (https://vimeo.com/1108446977); Low-Awareness + Centrality (https://vimeo.com/1108446940); High-Awareness + Efficient (https://vimeo.com/1108446899); High-Awareness + Centrality (https://vimeo.com/1108446923).

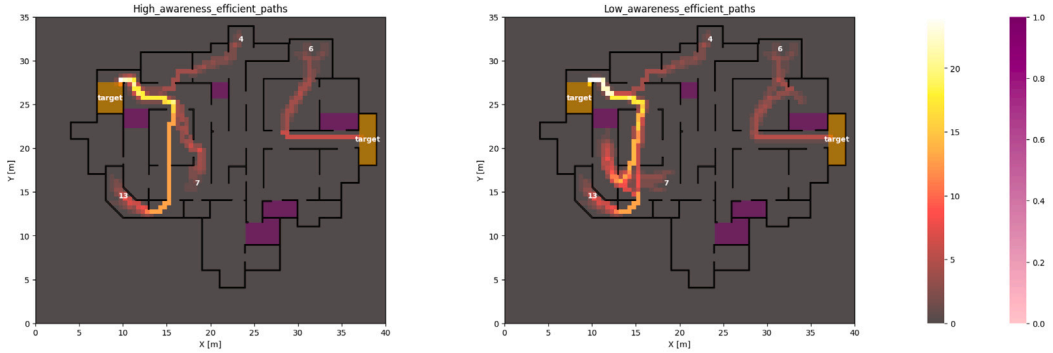

(a) High-Awareness – Efficient (b) Low-Awareness – Efficient

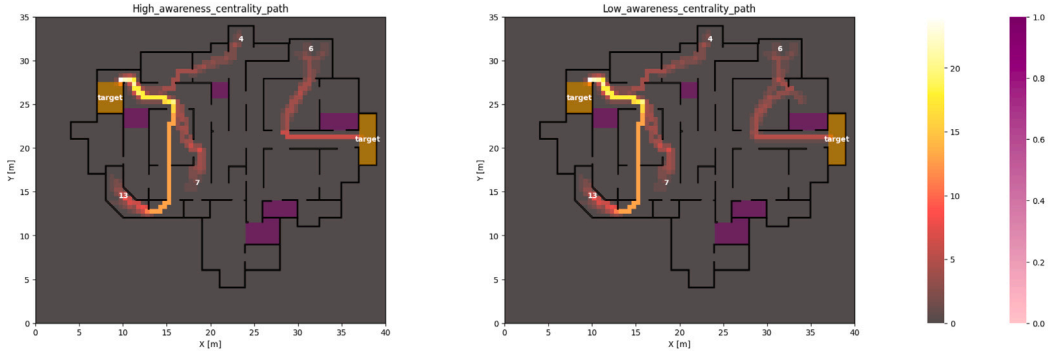

(c) High-Awareness – Centrality (d) Low-Awareness – Centrality

Fig. 6. Heatmaps showing unique agents per cell and the initial hazard sources for the representative Theme Park scenario under all routing and awareness configurations.

underlying spatial layout remained fixed, while variability was introduced through different initial agent distributions and hazard configurations. Complete recordings for the representative case are available.[5]

Fig. 7 shows snapshots of the risk field and agent distribution for the representative Cruise Ship scenario under the different routing and awareness configurations. These visualizations provided qualitative insight into how local risk perception and spatial constraints influenced agent behavior in a highly constrained environment.

To complement these qualitative observations, Table 6 reports aggregated performance metrics in terms of mean ± standard deviation computed over all Cruise Ship scenarios and simulation runs.

The aggregated results confirmed that situational awareness played a dominant role in constrained environments. Under low-awareness conditions, efficient-path routing exhibited the highest evacuation times and Mean RPR, together with substantially larger variability. In contrast, agile routing reduced both Mean RPR and performance variability, indicating more robust behavior than efficient-path routing when hazard information was delayed.

Importantly, this improvement was not achieved at the expense of substantially longer routes or evacuation times. As shown in Tables 5 and 6, agile routing yielded consistently lower Mean RPR in both environments, while evacuation times and remaining-path costs remained comparable to those obtained with efficient routing, and in several low-awareness settings even improved. Mean RPR indicated that agile routing was associated with lower overall risk exposure: by favoring structurally robust routes, it reduced the expected exposure along the remaining planned path without introducing a material time or path-length penalty.

Under high-awareness conditions, differences between routing strategies were reduced. Efficient routing achieved slightly lower evacuation times, while agile routing yielded marginally lower Mean RPR. These results suggested that, in highly constrained layouts, improved anticipatory hazard perception partially compensates for the structural robustness advantages of centrality-based routing.

Across both complex environments, the results consistently showed that higher situational awareness led to safer and more stable evacuation outcomes. At the same time, the advantages of agile routing became most pronounced under low-awareness conditions, particularly in spatially constrained environments, motivating the comparative analysis presented in the following section.

[5] Cruise Ship simulations: Low-Awareness + Efficient (https://vimeo.com/1105055814); Low-Awareness + Centrality (https://vimeo.com/1105057698); High-Awareness + Efficient (https://vimeo.com/1105055798); High-Awareness + Centrality (https://vimeo.com/1105057717).

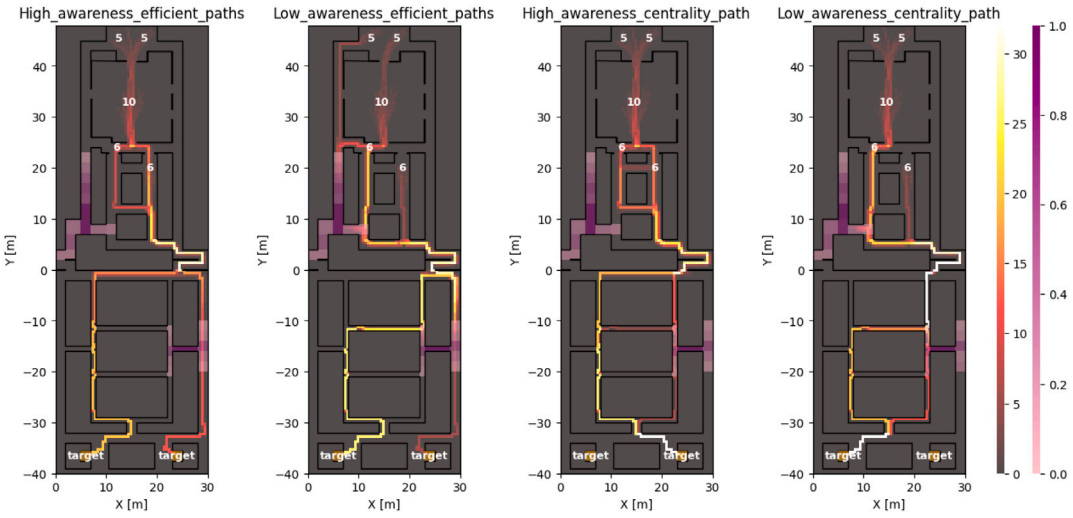

(a) High-Awareness – Efficient (b) Low-Awareness – Efficient (c) High-Awareness – Centrality (d) Low-Awareness – Centrality

Fig. 7. Snapshot of agent distribution and risk field for the representative Cruise Ship scenario under all routing and awareness configurations.

Table 6
Mean ± standard deviation of performance metrics aggregated over all Cruise Ship scenarios and simulation runs.

Configuration	Max Evac. Time (s)	Avg Evac. Time (s)	Avg remaining-path cost (s)	Mean RPR	RPR Var.
Centrality, High-Awareness	107.48 ± 30.78	101.61 ± 29.52	50.82 ± 19.78	0.0010 ± 0.0011	0.0002 ± 0.0003
Centrality, Low-Awareness	110.15 ± 33.89	105.79 ± 33.64	51.06 ± 21.38	0.0090 ± 0.0125	0.0034 ± 0.0053
Efficient, High-Awareness	102.89 ± 32.99	97.41 ± 31.63	47.33 ± 19.78	0.0010 ± 0.0019	0.0003 ± 0.0006
Efficient, Low-Awareness	118.20 ± 46.56	114.10 ± 46.18	47.73 ± 23.84	0.0183 ± 0.0214	0.0075 ± 0.0098

6.3. Comparative evaluation

Building on the environment-level results presented in Section 6.2, we now perform a systematic comparative analysis of the two routing strategies (Centrality vs. Efficient) under both situational awareness levels. In contrast to the aggregated results discussed previously, this section focuses on a case-wise comparison across the full set of simulated scenarios reported Appendix A. For each scenario, the two strategies are compared directly in terms of evacuation time and mean remaining-path risk (Mean RPR), allowing us to assess the consistency and robustness of their relative performance.

We first consider the Corridor environment (Fig. 2), which served as a baseline due to its simple topology and limited route diversity. In this environment, both routing strategies yielded equal outcomes according to the 3% criterion across all ten scenarios, as near-optimal paths largely coincided and hazard evolution offered little opportunity for structurally distinct rerouting. Consequently, neither evacuation time nor Mean RPR was meaningfully affected by the choice of routing heuristic.

Fig. 8 provides a compact overview of the three-way comparative outcomes (Centrality better, Efficient better, or equal performance) for the Theme Park and Cruise Ship environments, reported separately for evacuation time and Mean RPR. The Corridor environment yielded only equal outcomes across scenarios (as discussed above) and was therefore omitted from this summary. Outcomes were classified as *equal performance* when the relative difference between strategies was within 3%; otherwise, the strategy with the lower value was labeled as better.

In the Cruise Ship environment, agile routing exhibited a clear advantage under low-awareness conditions. For evacuation time (Fig. 8(a)), Centrality outperformed Efficient in the majority of low-awareness scenarios (10 cases versus 2), with no cases resulting in equal performance. A similar pattern was observed for Mean RPR (Fig. 8(b)), where Centrality dominated under low-awareness in 9 cases, while Efficient outperformed Centrality in only 2 cases and a single tie was observed. These results indicated that in spatially constrained layouts characterized by narrow corridors, compartmentalization, and structural bottlenecks, agile routing effectively mitigated early commitment to fragile paths and reduced the likelihood of high-risk rerouting cascades when hazard perception is delayed.

Under high-awareness conditions in the Cruise Ship environment, the difference between routing strategies became less pronounced. As illustrated in Fig. 8(a), Efficient routing more frequently achieved shorter evacuation times (8 cases versus 4), while the corresponding Mean RPR results (Fig. 8(b)) were more balanced, with Efficient outperforming Centrality in 6 cases, Centrality performing better in 5 cases, and a single case resulting in equal performance. This suggested that once agents could

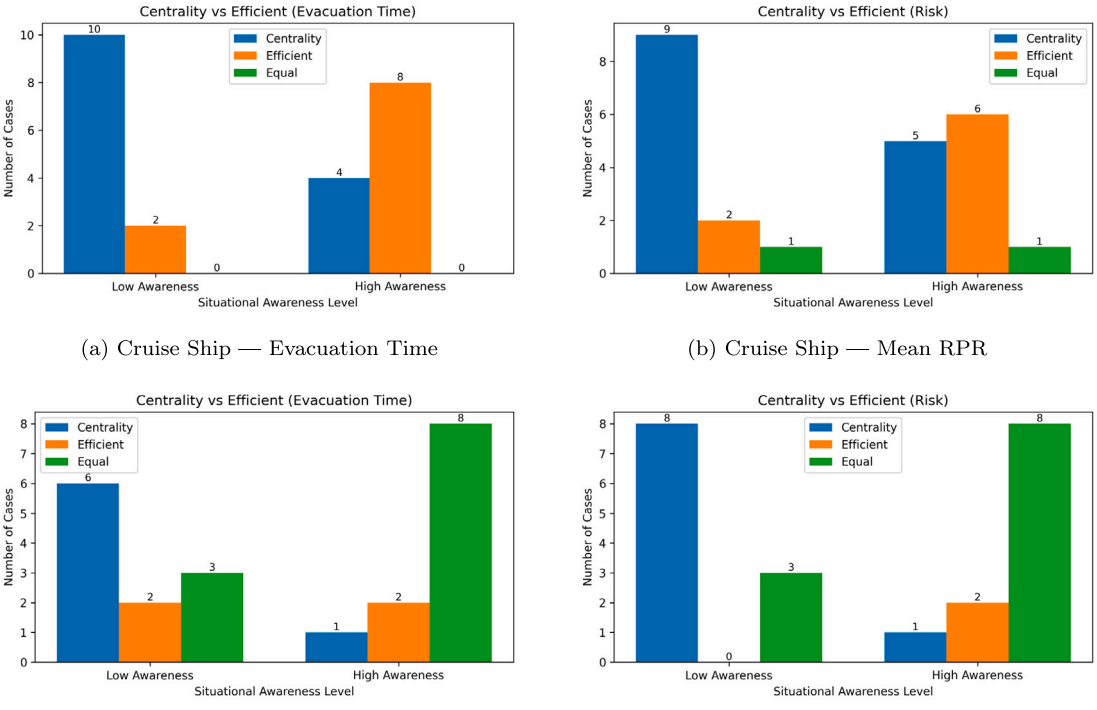

(a) Cruise Ship — Evacuation Time

(b) Cruise Ship — Mean RPR

(c) Theme Park — Evacuation Time

(d) Theme Park — Mean RPR

Fig. 8. Three-way comparative outcomes (Centrality better, Efficient better, or equal performance) for evacuation time and mean remaining-path risk across the Theme Park and Cruise Ship environments. Results are reported on a case-wise basis over all simulated scenarios.

reliably perceive evolving hazards, local responsiveness could partially compensate for the structural robustness advantages of centrality-based routing.

In the Theme Park environment, results were more mixed due to the high redundancy provided by wide corridors and the abundance of alternative routes. For evacuation time (Fig. 8(c)), Centrality performed better in half of the low-awareness scenarios (6 cases), while Efficient outperformed Centrality in fewer cases (2), with several ties (3), indicating largely comparable performance. Under high-awareness, most scenarios resulted in equal evacuation times (8 cases), with Efficient slightly outperforming Centrality more often than the reverse (2 versus 1 cases).

A similar trend was observed for Mean RPR in the Theme Park layout (Fig. 8(d)). Under low-awareness, Centrality showed a clearer advantage, outperforming Efficient in 8 scenarios, with no cases favoring Efficient and a small number of ties. Under high-awareness, however, the majority of scenarios again resulted in equal performance (8 cases), with Centrality and Efficient each outperforming the other in only a few cases (1 and 2 scenarios, respectively). Overall, these results suggested that in highly redundant environments the influence of the routing heuristic is reduced, and evacuation performance is increasingly governed by situational awareness and local hazard perception rather than global path structure.

Taken together, the case-wise comparative analysis confirmed that situational awareness was the dominant factor shaping evacuation outcomes across all considered environments. Agile routing proved particularly effective under low-awareness conditions, where it delayed commitment to structurally fragile paths, thereby improving robustness when hazard perception was delayed. Notably, these gains were primarily reflected in Mean RPR, which provided the clearest indicator of lower risk exposure, and they generally occurred without a corresponding increase in evacuation time or remaining-path cost.

6.4. On agent density, exit availability, and congestion effects

To complement the controlled experiments discussed above, an auxiliary simulation was conducted to examine whether agile routing led to congestion in structurally central areas under more realistic conditions. This scenario increased agent density and distributed multiple evacuation targets throughout the environment. Unlike previous setups, no dynamic hazards were introduced, allowing the analysis to focus exclusively on the interaction between routing logic, space usage, and congestion.

The cruise ship environment was reused for this experiment, with four emergency exits distributed across the layout to better reflect typical architectural redundancy and reduce artificial funneling. All agents followed agile routing under low situational awareness, allowing the simulation to probe whether the strategy's preference for topologically central paths resulted in excessive aggregation when risk was not actively shaping decisions. The full set of trajectories followed by the 96 agents is shown in Fig. 9,

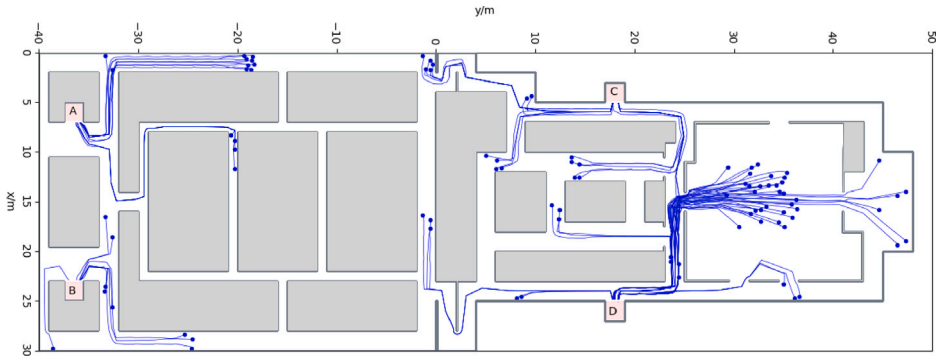

Fig. 9. Agent trajectory traces for the high-density cruise ship evacuation scenario (96 agents - please notice that the environment is rotated for better space utilization) using agile routing with low situational awareness. The plot depicts the paths followed by agents across the full evacuation process, with emergency exits labeled A, B, C, and D, shown with a light pink background and black lettering.

revealing the global flow structure and exit usage patterns. Complementarily, the temporal evolution of spatial speed distributions was captured in Fig. 10, which presents six representative snapshots illustrating the development and dissipation of local congestion zones.

Despite the increased number of agents, the speed profiles in Fig. 10 did not exhibit persistent or extreme congestion at any single location. This outcome could be attributed to the availability of multiple exits, which allowed agents to disperse across structurally distinct routes. Even though agile routing favors well-connected nodes, the resulting flow distribution remained balanced, and localized slowdowns dissipated as agents redistributed spatially.

These findings underscored the importance of environmental structure in shaping congestion outcomes. While centrality-based strategies could, in principle, concentrate flows around specific areas, their impact was highly dependent on exit placement and topological redundancy. In realistic layouts with multiple accessible targets, the strategy did not inherently produce bottlenecks.

Nonetheless, in large-scale evacuations with very high densities or with hazards emerging near central hubs, congestion effects could become more pronounced. Addressing such scenarios would require routing mechanisms that explicitly account for crowding dynamics. The integration of congestion-aware path planning into the current framework represents a promising avenue for future work.

7. Conclusion

This paper presented a simulation-based comparison of two routing strategies: an efficiency-driven approach aimed at minimizing nominal evacuation time, and an agile routing strategy that prioritizes structural adaptability by maximizing evacuation betweenness centrality within the γ-efficient path set, under varying levels of situational awareness. Here, γ denotes a tolerance parameter that defines the set of admissible routes as those whose travel cost does not exceed $(1 + \gamma)$ times the shortest-path cost; in this sense, agents are allowed to consider not only the strictly optimal path but also near-optimal alternatives within a bounded efficiency margin. Through the development of a modular simulation framework extending JuPedSim and a carefully designed set of controlled experiments, we evaluated how autonomous agents respond to dynamically evolving hazards under different combinations of routing strategies and situational awareness levels, and how these factors comparatively influence evacuation efficiency and risk exposure.

Beyond the specific results obtained, this work provides a structured baseline for exploring the interaction between risk awareness and routing logic in dynamic environments. The proposed framework combines microscopic physical motion modeling with awareness-driven routing decisions within a modular architecture, contributing to the growing line of research on hybrid evacuation models that integrate perception, communication, and routing mechanisms [9,15,16]. The simulated scenarios were intentionally simplified to allow controlled comparisons between routing strategies and awareness levels, while still capturing non-trivial dynamics that mirror recent findings in empirical and virtual-reality studies on situational awareness and guidance [2,3].

Two key findings emerged from the study. First, situational awareness plays a decisive role in evacuation performance. Across all scenarios, agents equipped with higher awareness — capable of evaluating evolving risks across multiple nodes — consistently outperformed their less-informed counterparts. These agents demonstrated earlier rerouting, better hazard avoidance, and fewer last-minute deviations, resulting in faster and safer evacuations. This confirms and extends prior evidence that real-time risk information, when effectively perceived and communicated, substantially improves collective evacuation outcomes [2,9].

Second, agile routing exhibited distinct advantages in scenarios characterized by complex environmental structures, especially when some alternative evacuation routes became unavailable due to increasing risk levels. While efficient-path routing, based on minimizing nominal travel time, was adequate in static or predictable environments, it may direct agents toward routes that later become unsafe under dynamic hazard conditions. By contrast, centrality-aware agents naturally favored topologically significant paths, namely those more likely to preserve access to alternative safe routes as conditions evolved. Although centrality does not guarantee the minimum nominal travel time, it enhances flexibility and connectivity, two properties that proved especially valuable

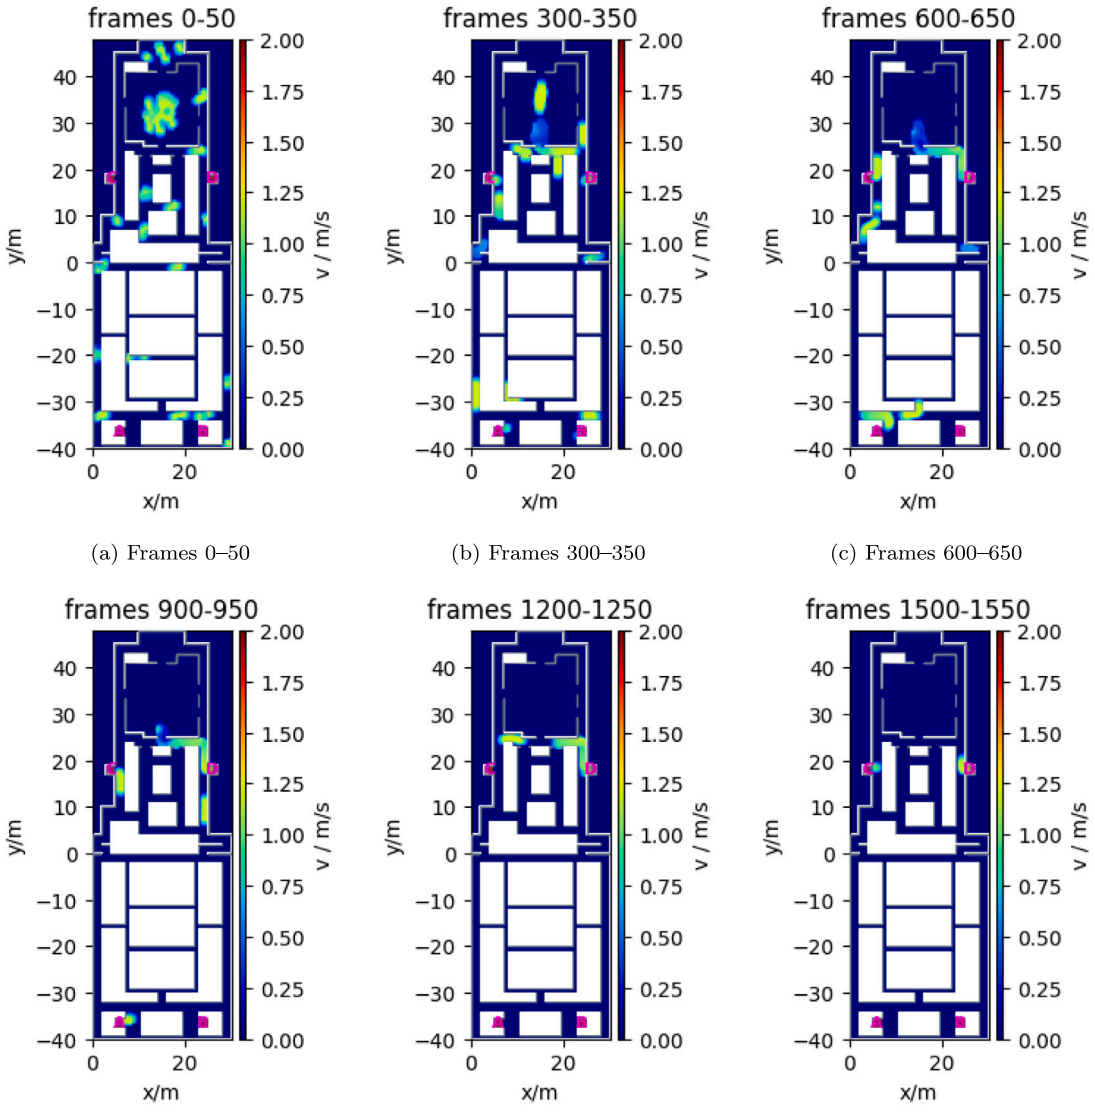

(a) Frames 0–50 (b) Frames 300–350 (c) Frames 600–650

(d) Frames 900–950 (e) Frames 1200–1250 (f) Frames 1500–1550

Fig. 10. Spatial speed profiles for a high-density cruise ship evacuation scenario (96 agents) using agile routing with low situational awareness at different temporal stages. Emergency exits are annotated with magenta labels corresponding to Exits A, B, C, and D.

when detours or route revisions became necessary. These results indicate that evacuation robustness benefits from structurally adaptable routing choices, particularly when hazard information is limited or delayed [14,16].

The effectiveness of agile routing is inherently dependent on the topological structure of the evacuation environment, and in particular on the distribution and accessibility of exit nodes. In layouts characterized by high structural redundancy and multiple alternative routes, agile routing can exploit network connectivity to favor structurally adaptable paths, thereby reducing risk exposure. However, in environments with limited route diversity, such as single-channel architectures, linear corridors, or configurations where paths converge toward a single bottleneck, the advantages of agile routing are significantly reduced. In such cases, near-optimal paths tend to overlap structurally, and routing decisions become largely constrained by the underlying topology rather than by the selected routing strategy. This limitation is consistent with the results observed in the Corridor environment, where both routing strategies produced equivalent outcomes. These findings highlight that the benefits of agile routing should be interpreted in conjunction with the spatial design of evacuation infrastructures.

Taken together, these findings suggest that effective evacuation is not simply a matter of minimizing nominal travel time, but of balancing efficiency, situational awareness, and structural adaptability. In uncertain contexts where threats evolve during the evacuation process, agents benefit from paths that offer structural options and continuous information updates. In practice,

this implies that smart infrastructures and support systems should not only provide direction but also awareness—through visual, auditory, or digital cues such as adaptive signage, mobile alerts, or ambient feedback systems [30,31]. The integration of perception-aware technologies into evacuation management may represent a key step toward more resilient crowd control and safety planning.

Nonetheless, the present study has several limitations. To keep the analysis focused and reproducible, the model abstracts away psychological and social processes (e.g., panic reactions, trust dynamics, and communication breakdowns) and assumes compliance with guidance recommendations, whereas future evaluations should examine how partial compliance or heterogeneous behavioral responses affect the observed trends. Situational awareness is represented using a two-level abstraction (high vs. low) to enable controlled comparisons, although real-world awareness is likely to vary continuously and stochastically as a function of perception and communication constraints; extending the model toward graded or uncertainty-aware awareness representations is therefore a relevant direction. Finally, hazard propagation is modeled deterministically to facilitate attribution of effects across scenarios, at the expense of explicitly representing uncertainty and noise, suggesting the value of complementary analyses that assess sensitivity to alternative propagation assumptions; in this context, the proposed group behavior mechanism offers a structured starting point for representing group-level effects and can be further evaluated across a broader range of group configurations.

The presented approach takes a traditional perspective both on pedestrian modeling (we essentially leverage prior work through the open-source JuPedSim simulator, only adding a concept of group that influences path planning, and that will require further investigation) as well as on risk simulation (defined and implemented through simplified hand–crafted propagation rules) and pedestrian routing, based on different strategies (agile and efficient path selection). While the routing module takes an optimization perspective, not exactly requiring the definition of rules but rather allowing the expression of a mathematical formulation of the structure of the environment, routing dynamics, the concept of effectiveness and efficiency of evacuation, the overall approach is completely *white–box* and *explainable*: overall properties for routes can be enforced and verified. Pedestrian models obtained as results of ML approaches could, in principle, be integrated as substitutes of existing operational level models existing in JuPedSim, but – besides being outside of the scope of the present work – it would imply concerns on the explainability of the achieved results. The widespread and impressive success of Large Language Models (LLMs) as flexible components that can be employed in a wide range of scientific activities raises the question whether these tools could be applied to our case. Instead of using an LLM as an *alternative* to optimization approaches, we could employ them as an orchestrator analyzing the overall state of affairs in the environment (e.g. presence of hazards and their trend, current position of pedestrians and their trajectories) to investigate options proposed by the different routing strategies, evaluate them, potentially propose integration and/or modifications. Our position is that, however, such a system would – as of this moment – require a form of human oversight, due to ethical issues and very practical ones (i.e. explainability, errors or confabulations). This might be an interesting but still exploratory line for future works, especially should novel results on trustworthy usages of agentic AI support real-world usage of such systems.

Future work will explore extensions of the framework toward probabilistic and perception-aware representations of situational awareness, investigate congestion-aware routing, and assess behavior under higher-density scenarios. Where suitable data and experimental settings are available, controlled studies (e.g., virtual-reality setups) and sensor-informed scenarios may help constrain model parameters and characterize decision latency. Overall, these directions aim to move toward evacuation support systems that incorporate richer information while explicitly managing the trade-offs between safety, efficiency, and human decision-making.

Code availability

Source code and experiment scripts are available at: https://github.com/AlvaroS3rrano/Evacuation_Simulation/.

Acknowledgments

Álvaro Serrano and Marin Lujak were supported by the VAE grant: TED2021-131295B-C33 funded by MCIN/AEI/10.13039/501100 011033 and by the European Union's "NextGenerationEU/PRTR", under "ERDF A way of making Europe", by the grant COSASS: PID2021-123673OB-C32 funded by MCIN/AEI/10.13039/501100011033 and by the EVASAI grant: PID2024-158227NB-C32 funded by MICIU/AEI/10.13039/501100011033/FEDER, UE. Giuseppe Vizzari was supported by the Spoke 8 — MaaS and Innovative services of the National Center for Sustainable Mobility (MOST) set up by the "Piano nazionale di ripresa e resilienza (PNRR)" — M4C2, founding 1.4, funded by the European Union (Project code CN00000023, CUP D93C22000410001), and by the MUR under the grant "Dipartimenti di Eccellenza 2023-2027" of the DISCO Department University of Milano-Bicocca, Italy.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.simpat.2026.103294.

Data availability

No data was used for the research described in the article.

References

[1] C. Feliciani, A. Corbetta, M. Haghani, K. Nishinari, Trends in crowd accidents based on an analysis of press reports, Saf. Sci. 164 (2023) 106174, http://dx.doi.org/10.1016/j.ssci.2023.106174.

[2] P. Zhang, L. Yang, S. Lo, D. Wang, M. Li, J. Jiang, N. Jiang, Experimental study on evacuation behavior with guidance under high and low urgency conditions, Saf. Sci. 154 (2022) 105865, http://dx.doi.org/10.1016/j.ssci.2022.105865.

[3] J. Wang, L. Shen, H. Wang, R. Lovreglio, Y. Tong, An empirical investigation on the impact of evacuation signages on individual behaviour in building evacuations, J. Build. Eng. 111 (2025) 113519, http://dx.doi.org/10.1016/j.jobe.2025.113519.

[4] A. Mossberg, D. Nilsson, H. Frantzich, Evaluating new evacuation systems related to human behaviour using a situational awareness approach, Fire Saf. J. 134 (2022) 103693, http://dx.doi.org/10.1016/j.firesaf.2022.103693.

[5] D. Zhu, H. Zhou, A. Lin, B. Zhou, A dynamic risk-aware routing recommendation using deep reinforcement learning, in: Proceedings of the IEEE International Conferences on Internet of Things, Green Computing and Communications, Cyber, Physical and Social Computing, Smart Data and Congress on Cybermatics, IEEE, 2024, pp. 235–240, http://dx.doi.org/10.1109/ITHINGS-GREENCOM-CPSCOM-SMARTDATA-CYBERMATICS62450.2024.00058.

[6] Y. Tong, N.W.F. Bode, M. Haghani, R. Lovreglio, Exploring occupant exit choices during fire drills and false alarm evacuations in a library, Saf. Sci. 182 (2025) 106708, http://dx.doi.org/10.1016/j.ssci.2024.106708.

[7] G. Vizzari, L. Crociani, S. Bandini, An agent-based model for plausible wayfinding in pedestrian simulation, Eng. Appl. Artif. Intell. 87 (2020) 103241, http://dx.doi.org/10.1016/j.engappai.2019.103241.

[8] A.U. Kemloh Wagoum, A. Seyfried, S. Holl, Modeling the dynamic route choice of pedestrians to assess the criticality of building evacuation, Adv. Complex Syst. 15 (7) (2012) 1250029, http://dx.doi.org/10.1142/S0219525912500294.

[9] M. Keykhaei, N. Neysani Samany, M. Jelokhani-Niaraki, S. Zlatanova, Multi-agent-based human cognition simulation of situation-aware earthquake emergency evacuation, Int. J. Disaster Risk Reduct. 100 (2024) 104183, http://dx.doi.org/10.1016/j.ijdrr.2023.104183.

[10] A. Zablotsky, M. Kuperman, S. Bouzat, Pedestrian evacuations with imitation of cooperative behavior, Phys. Rev. E 109 (2024) 054304, http://dx.doi.org/10.1103/PhysRevE.109.054304.

[11] A.U. Kemloh Wagoum, B. Steffen, A. Seyfried, M. Chraibi, Parallel real-time computation of large-scale pedestrian evacuations, Adv. Eng. Softw. 60–61 (2013) 98–103, http://dx.doi.org/10.1016/j.advengsoft.2012.10.001.

[12] J. Wąs, R. Lubaś, Towards realistic and effective agent-based models of crowd dynamics, Neurocomputing 146 (2014) 199–209, http://dx.doi.org/10.1016/j.neucom.2014.04.057.

[13] M. Lujak, S. Ossowski, Evacuation route optimization architecture considering human factor, AI Commun. 45 (2017) 78–89, http://dx.doi.org/10.3233/AIC-170721.

[14] M. Lujak, S. Giordani, Centrality measures for evacuation: Finding agile evacuation routes, Future Gener. Comput. Syst. 83 (2018) 401–412, http://dx.doi.org/10.1016/j.future.2017.05.014.

[15] P. Zhang, W. Liu, L. Yang, J. Wu, K. Wang, Y. Cui, A social force-based model for pedestrian evacuation with static guidance in emergency situations, Fire 8 (1) (2025) 30, http://dx.doi.org/10.3390/fire8010030.

[16] X. Liu, T. Zhou, H. Kang, J. Ma, Z. Wang, J. Huang, W. Weng, Y.-K. Lai, K. Li, RESCUE: Crowd evacuation simulation via controlling SDM-united characters, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV, 2025, pp. 24955–24964, URL https://arxiv.org/abs/2507.20117.

[17] C.M. Mayr, A. Templeton, G. Köster, Designing mobile application messages to impact route choice: A survey and simulation study, PLoS One 18 (4) (2023) e0284540, http://dx.doi.org/10.1371/journal.pone.0284540.

[18] C.M. Mayr, G. Köster, Guiding crowds when facing limited compliance: Simulating strategies, PLoS One 17 (11) (2022) e0276229, http://dx.doi.org/10.1371/journal.pone.0276229.

[19] JuPedSim Contributors, JuPedSim documentation, 2024, URL https://www.jupedsim.org. (Accessed 30 May 2025).

[20] W.W.Y. Lam, S. Yao, B.P.Y. Loo, Pedestrian exposure measures: A time-space framework, Travel. Behav. Soc. 1 (1) (2014) 22–30, http://dx.doi.org/10.1016/j.tbs.2013.10.004.

[21] D. Helbing, P. Molnár, Social force model for pedestrian dynamics, Phys. Rev. E 51 (5) (1995) 4282–4286, http://dx.doi.org/10.1103/PhysRevE.51.4282.

[22] D. Helbing, I. Farkas, T. Vicsek, Simulating dynamical features of escape panic, Nature 407 (6803) (2000) 487–490, http://dx.doi.org/10.1038/35035023.

[23] M. Chraibi, A. Seyfried, A. Schadschneider, Generalized centrifugal-force model for pedestrian dynamics, Phys. Rev. E 82 (2010) 046111, http://dx.doi.org/10.1103/PhysRevE.82.046111, URL https://link.aps.org/doi/10.1103/PhysRevE.82.046111.

[24] C. Burstedde, K. Klauck, A. Schadschneider, J. Zittartz, Simulation of pedestrian dynamics using a two-dimensional cellular automaton, Phys. A 295 (3) (2001) 507–525, http://dx.doi.org/10.1016/S0378-4371(01)00141-8, URL https://www.sciencedirect.com/science/article/pii/S0378437101001418.

[25] A. Tordeux, M. Chraibi, A. Seyfried, A. Schadschneider, Prediction of pedestrian dynamics in complex architectures with artificial neural networks, J. Intell. Transp. Syst. 24 (6) (2020) 556–568, http://dx.doi.org/10.1080/15472450.2019.1621756.

[26] G. Vizzari, T. Cecconello, Pedestrian simulation with reinforcement learning: A curriculum-based approach, Futur. Internet 15 (1) (2023) http://dx.doi.org/10.3390/fi15010012, URL https://www.mdpi.com/1999-5903/15/1/12.

[27] A. Schadschneider, W. Klingsch, H. Klüpfel, T. Kretz, C. Rogsch, A. Seyfried, Evacuation Dynamics: Empirical Results, Modeling and Applications, Springer, Berlin, 2010.

[28] A. Abdelghany, K. Abdelghany, H. Mahmassani, W. Alhalabi, Modeling framework for optimal evacuation of large-scale crowded pedestrian facilities, J. Adv. Transp. 48 (6) (2014) 694–711, http://dx.doi.org/10.1002/atr.1260.

[29] E. Andresen, M. Chraibi, A. Seyfried, A representation of partial spatial knowledge: a cognitive map approach for evacuation simulations, Transp. A: Transp. Sci. 14 (5–6) (2018) 433–467, http://dx.doi.org/10.1080/23249935.2018.1432717.

[30] A. Corbetta, T. Senan, L. Wöstemeier, B. Hengeveld, Public-space sonification for pedestrian trajectory nudging, in: K.R. Rao, A. Seyfried, A. Schadschneider (Eds.), Traffic and Granular Flow '22, Springer Nature Singapore, Singapore, 2024, pp. 207–214, http://dx.doi.org/10.1007/978-981-99-7976-9_26.

[31] C. Feliciani, K. Shimura, K. Nishinari, Crowd control methods: Established and future practices, in: Introduction To Crowd Management: Managing Crowds in the Digital Era: Theory and Practice, Springer International Publishing, Cham, 2021, pp. 167–216, http://dx.doi.org/10.1007/978-3-030-90012-0_6.

[32] R.-Y. Guo, H.-J. Huang, S.C. Wong, Route choice in pedestrian evacuation under conditions of good and zero visibility: Experimental and simulation results, Transp. Res. Part B: Methodol. 46 (6) (2012) 669–686, http://dx.doi.org/10.1016/j.trb.2012.01.002.

[33] D. Helbing, P. Mukerji, Crowd disasters as systemic failures: Analysis of the Love parade disaster, EPJ Data Sci. 1 (1) (2012) 1–40, http://dx.doi.org/10.1140/epjds7.

[34] A. Sieben, A. Seyfried, Inside a life-threatening crowd: Analysis of the Love parade disaster from the perspective of eyewitnesses, Saf. Sci. 166 (2023) 106229, http://dx.doi.org/10.1016/j.ssci.2023.106229.

[35] J.Y. Yen, Finding the k shortest loopless paths in a network, Manag. Sci. 17 (11) (1971) 712–716, http://dx.doi.org/10.1287/mnsc.17.11.712.

[36] J. Hershberger, M. Maxel, S. Suri, Finding the k shortest simple paths: A new algorithm and its implementation, ACM Trans. Algorithms (TALG) 3 (4) (2007) 45–es.

[37] R. Lovreglio, E. Ronchi, D. Nilsson, A model of the decision-making process during pre-evacuation, Fire Saf. J. 78 (2015) 168–179, http://dx.doi.org/10.1016/j.firesaf.2015.07.001, URL https://www.sciencedirect.com/science/article/pii/S0379711215300023.

[38] F. Zanlungo, Z. Yücel, D. Brščić, T. Kanda, N. Hagita, Intrinsic group behaviour: Dependence of pedestrian dyad dynamics on principal social and personal features, PLoS One 12 (11) (2017) 1–26, http://dx.doi.org/10.1371/journal.pone.0187253.

[39] S. Bandini, A. Gorrini, G. Vizzari, Towards an integrated approach to crowd analysis and crowd synthesis: A case study and first results, Pattern Recognit. Lett. 44 (2014) 16–29, http://dx.doi.org/10.1016/j.patrec.2013.10.003, Pattern Recognition and Crowd Analysis, URL https://www.sciencedirect.com/science/article/pii/S0167865513003760.

[40] C. Feliciani, X. Jia, H. Murakami, K. Ohtsuka, G. Vizzari, K. Nishinari, Social groups in pedestrian crowds as physical and cognitive entities: Extent of modeling and motion prediction, Transp. Res. Part A: Policy Pr. 176 (2023) 103820, http://dx.doi.org/10.1016/j.tra.2023.103820, URL https://www.sciencedirect.com/science/article/pii/S0965856423002409.

[41] D. Weyns, A. Omicini, J. Odell, Environment as a first class abstraction in multiagent systems, Auton. Agents Multi-Agent Syst. 14 (1) (2007) 5–30, http://dx.doi.org/10.1007/s10458-006-0012-0.