

# The Cost Of The Omniscience Assumption: Quantifying Bias In Fire Evacuation Predictions Under Incomplete Information

Desheng Cao<sup>1</sup>, Xin Liu<sup>1</sup>, Jingjing Xia<sup>1</sup>, Qingxin Xia<sup>1,\*</sup>

<sup>1</sup> School of Computer and Information Security, University of Emergency Management, Beijing, China

| ARTICLE INFO                                                                                                                                                                                                                                                                                                                                      | ABSTRACT                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Article history:</b><br/>Received 23 June 2026<br/>Received in revised form 17 August 2026<br/>Revised: [to be completed by the journal]<br/>Accepted 1 October 2026<br/>Available online 5 October 2026</p> <p><b>Keywords:</b><br/>Omniscience assumption; Occupant evacuation; Prediction bias; Multi-agent systems; Fire simulation</p> | <p>Evacuation simulation is the main basis for determining safety margins in performance-based fire design, and its reliability depends on assumptions about the state of crowd information. Existing microscopic evacuation models generally imply an omniscient assumption: individuals know the locations of all exits and fire sources at any time, and information is received instantly. This contradicts behavioral evidence from real fires, where perception is limited, cognition is insufficient, and information is delayed. The prediction errors caused by this assumption haven't been systematically quantified yet. This paper breaks down the omniscience assumption into three orthogonal information channels: perception, cognition, and delay, represented respectively by the visual radius that diminishes with smoke density, the completeness of knowledge about exits and fire sources, and the update interval for decision information. On a multi-agent platform coupling fire, smoke, and humans, we designed gradient ablation experiments: first downgrading each channel separately to isolate their individual contributions, then combining downgrades to identify channel interactions. Deviations are measured across four aspects: evacuation time T95, number of casualties, safety margin distribution, and balance in exit usage. Experiments show that the cognitive channel is the most stable source of independent bias (dominating 16 out of 18 scenarios), while the perception and delay channels only appear under high-density conditions with few exits and smoke spread, and the channels work together only when there's a cognitive gap. This suggests that an all-knowing model is a pretty good approximation for normal real-world scenarios, and its cost is mainly in the extreme cases where information is completely missing, providing a quantitative basis for prioritizing prediction corrections and information system investments.</p> |

## 1. Introduction

In performance-based fire design, whether a building design is safe depends on comparing the required safe egress time (RSET) with the available safe egress time (ASET), where RSET is given by an evacuation simulation model. Reviews by Gwynne and others show that evacuation modeling has developed over the years into three main approaches: optimization, simulation, and risk assessment.

\* Corresponding author.  
E-mail address: [chenggang@ncist.edu.cn](mailto:chenggang@ncist.edu.cn)

Just the simulation models alone have dozens of implementations, and any model has to make assumptions about the building space, the composition of people, and crowd behavior [1]. On a microscopic level, cellular automata (CA) and floor field models describe individual movement and competition at exits using local transition rules. Because the rules are simple and highly parallelizable, they've become one of the most commonly used modeling paradigms [2, 3]. The model's outputs for evacuation time, congestion spots, and casualty risks directly feed into the design criteria. Any systematic deviation from real behavior would proportionally translate into misjudging safety margins: for example, if the model consistently overestimates people's ability to process information, the evacuation time it predicts will be systematically too short, and the design might be wrongly considered safe.

The reliability of behavioral assumptions thus becomes the bottleneck for the predictive capability of evacuation models. Wang and others, in a systematic review of human factors in emergency evacuations, pointed out that behavioral uncertainty is one of the main sources of prediction uncertainty [2]. Among various behavioral assumptions, this paper focuses on a type that is common across models but rarely tested on its own: individuals having perfect information. This isn't a choice of implementation for a specific model, but a default premise in classic constructs like floor fields and distance fields, and its predictive cost hasn't been quantified so far. The omniscient assumption means that a model assumes each individual has three types of information abilities: perceptually, all exits and hazards throughout the building are visible; cognitively, the locations of all exits, the building layout, and the fire status are known in advance; temporally, there's no delay in acquiring or updating environmental information. In classic CA models, individuals choose which way to move based on global static or dynamic fields, and the exit attraction field is equally visible to everyone [2, 3], implicitly covering the first two abilities.

This paper formalizes omniscience into three orthogonal information channels: perception, cognition, and delay, and defines a gradient degradation process from omniscience to scarcity, making it possible to 'turn off a hypothesis' as a repeatable experiment; single-channel and combined ablation tests are performed on a fire–smoke–people coupled simulation platform to quantify prediction bias and attribute it to channels; matrix experiments with 18 basic scenarios are used to characterize the conditions under which the dominant channel shifts.

## **2. Related Work**

### *2.1 Evacuation Simulation Models*

Evacuation simulation models can be divided into three types based on granularity: macroscopic, mesoscopic, and microscopic. Microscopic models are further divided into continuous space models (represented by social force models [3]) and discrete space models (represented by cellular automata [4]). CA models are widely used because of their simple rules and good parallelism. Kirchner and Schadschneider proposed a biomimetic CA evacuation model, which simulates congestion and queueing at exits through conflict resolution probabilities [5]. Burstedde and others proposed a 2D CA model that uses a diffusive and decaying floor field to achieve self-organization of crowds, and they specifically pointed out that this model can simulate evacuation in environments with reduced visibility [6]. These classic models form the foundation of evacuation simulations, and their common feature is that individuals perceive the environment globally and without noise: the floor field, distance field, and exit field are equally visible to all individuals, and there is no information loss in decision-making inputs.

Evacuation simulation in fire scenarios needs to capture the interaction between fire, smoke, and people at the same time. Fire hazard environments are usually described by visibility, temperature, radiation, and toxic gas concentration. The cumulative effects of toxicity and heat exposure are often

assessed using the Fractional Effective Dose (FED) model, where an individual's FED reaching 1 at any moment is considered incapacitation [7]. In terms of coupling methods, Zhai Yue and others built a dynamic coupled model for fire and mixed pedestrian-vehicle evacuation in underground parking lots based on a cellular automaton background field and PyroSim fire data. They found that the necessary safe evacuation time in mixed pedestrian-vehicle cases is 20% to more than twice as long as that of purely pedestrian evacuation [8]. Marlow and others coupled computational fluid dynamics (CFD) based on the lattice Boltzmann method with the social force model to simulate pollutant dispersion in evacuation scenarios, allowing environmental exposure parameters to be given for each individual during evacuation [9]. Mirahadi and colleagues proposed an IFC-driven performance-based evacuation assessment framework that combines fire dynamics simulation with agent-based modeling to support multi-scenario evaluation during the design stage [10]. Xie and others conducted coupled simulations of agent behavior and CFD for fires on subway platforms, pointing out that the interaction between fire and crowds cannot be ignored, and that the fire source location and crowd density significantly affect evacuation outcomes [11]. The indoor evacuation model developed by Cao and others quantifies the local fire risk individuals face along evacuation routes and looks at how evacuation performance changes at different stages of fire development [12], with fire scene simulations mainly represented by computational fluid dynamics tools like FDS [13].

Looking at this approach, the dynamics of fire scenes have been pushed down to the level of individual units, yet the way individuals get information about the fire still follows a global, real-time idealized setup. The gap between a dynamic environment and all-knowing individuals is exactly the opening this article dives into.

## *2.2 Modeling Incomplete Information and Bounded Rationality*

The route closest to this study explicitly introduces incomplete information into individual decision-making. Liu and others incorporated gas concentration gradients into exit choice models, taking into account individuals' perceived distance of danger and their trust in others' information. The results showed that increasing perceived distance or trust could improve individual safety, while higher crowd density weakens the effects of perception and trust, leading to more exit-switching behavior [14]. Wang and others proposed a bounded rationality route choice model, using two parameters—level of caution and congestion sensitivity—to describe how individuals balance distance and crowding. They found that different combinations of bounded rationality significantly change the equilibrium of exit usage, which in turn affects the ranking of importance for evacuation facilities [15]. Yen and others proposed algorithms for individual evacuation based on smoke perception and group evacuation based on congestion perception, incorporating real-time smoke toxicity information into route planning and evacuation scheduling [16]. Ji and colleagues combined an improved CA model with IoT systems to provide dynamic evacuation guidance based on real-time fire and crowd data, and experiments showed that optimal real-time guidance can significantly shorten evacuation time [17]. More recent progress has brought large language models (LLMs) into evacuation decision simulation: Mendoza and others developed a hierarchical generative agent framework to simulate sequential decision-making with personalized agents, partly because existing models assume rational and homogeneous behavior, leading to overly optimistic predictions [18].

Although behavioral uncertainty has long been recognized as a key issue in evacuation modeling [2, 19], the existing approaches actually sidestep the omniscient assumption itself: one approach adds more detailed behavioral mechanisms within the model, like perception modeling and limited rational route choice, making individuals behave more realistically in specific scenarios; the other estimates output ranges through multiple random runs, capturing the discreteness of results [19]. Neither of them answers the question, 'If we stick to the omniscient assumption, how much would

predictions be off and where would the deviations come from?' This article breaks this overall issue down into three research questions that can be tested: RQ1 is about the scale of bias, meaning compared to an all-knowing baseline, how much do incomplete perception, cognition, and delay cause deviations in evacuation time, casualties, and safety margins, and whether the direction of the bias is consistently overly optimistic; RQ2 is about the attribution of bias, that is, how much each of the three channels contributes, whether there are interactions between channels, and how the dominant channel shifts depending on the fire source location, exit configuration, and crowd density; RQ3 is about outcome translation, meaning how the quantified breakdown of biases can be used to correct predictions from the all-knowing model and guide priority investments in evacuation information systems.

### 3. Formalization of the Omniscience Assumption and Its Decomposition into Three Channels

#### 3.1 Definition of the Omniscience Assumption

Suppose the building space  $S$  contains a set of exits  $E$  and a set of fire sources  $F$ . The information state of individual  $i$  at time  $t$  is represented as a triple, as shown in Equation (1).

$$I(i, t) = (P(i, t), K(i), D(i, t)) \quad (1)$$

Here,  $P$  is perceptual information, referring to the environmental state that an individual can observe at the current moment;  $K$  is cognitive information, referring to the spatial and hazard knowledge that the individual already has, independent of time;  $D$  is temporal information, indicating how outdated the individual's information is relative to the current environmental state. The omniscient assumption is that the three elements take their optimal values:

$P = P^*$ : unrestricted visibility, without occlusion or noise;

$K = K^*$ : all exits and fire sources are known;

$D = 0$ : information is updated in real time.

The distance between a real individual's information state and the optimal value is basically how incomplete the information is. This incompleteness can happen in several ways, like wrong information or asymmetry between individuals. In this article, we only focus on three dimensions: perception, cognition, and delay; the other aspects are left for future research.

#### 3.2 The Perception, Cognition, And Delay Channels

The three channels are classified according to the information lifecycle: perception determines what can be seen right now, cognition determines what is known beforehand, and delay determines how often information is updated. The three are orthogonal, meaning that downgrading one doesn't change the meaning of the other two, which provides independent variables for ablation experiments.

The perception channel  $P$  is characterized by an individual's visible radius  $r_v$ , which decreases with the local smoke concentration  $C$ , as shown in equation (2). There is experimental evidence that smoke visibility affects evacuation speed [20], and Isobe and others' zero-visibility room evacuation experiments further confirmed that loss of visibility has a significant impact on exit choice and evacuation time [21].

$$r_v = R_0 \cdot \exp(-\alpha C) \quad (2)$$

Here,  $R_0$  is the maximum visible radius when there's no smoke, and  $\alpha$  is the attenuation coefficient. The perception channel has three levels: the omniscient level with an infinite  $r_v$ , which is like global awareness; the partial level, where  $r_v$  has a finite value and varies with the local smoke concentration; and the scarce level, where  $r_v$  is almost just one cell, so individuals can only sense their

immediate surroundings, needing to follow walls or stick with others. The downgrade in perception directly affects whether someone can find exits, fire sources, or smoke boundaries in time.

The cognitive channel K is represented by the set of exits  $E_{\text{known}}$  that an individual knows and a flag indicating fire source knowledge. In the all-knowing case,  $E_{\text{known}} = E$  and the fire source location is known; in the partial case,  $E_{\text{known}}$  is a subset of exits familiar to the individual and the fire source location is unknown, so the individual prefers to head to familiar exits; in the scarce case,  $E_{\text{known}}$  is empty, so the individual can only choose a direction at random or follow nearby people. Cognitive degradation determines the quality of an individual's initial path choice, and previous studies have shown that exit selection behavior is significantly influenced by spatial familiarity and fire source knowledge [22].

The delay channel D is characterized by the decision information update interval  $\tau$ . The omniscient setting is  $\tau = 0$ , meaning individuals use the latest environment state at each decision step; the partial setting has  $\tau$  ranging from a few seconds to tens of seconds, during which individuals use the last captured snapshot of the environment; the scarce setting occurs when  $\tau$  approaches infinity, so individuals rely on the initial information from the start of evacuation throughout. Delay degradation reflects the lag in individuals' responses to environmental changes. The three channel settings correspond to the heterogeneity in real populations: the perception setting corresponds to differences in eyesight and field of view, the cognition setting to differences in familiarity with the building, and the delay setting to differences in reaction times. Therefore, ablation results can be directly mapped to modeling population heterogeneity.

### 3.3 Metrics Of Prediction Bias

Using the omniscient baseline as a reference, we define the prediction error for the ablation scenario. The main metric is the relative evacuation time deviation, as shown in Equation (3).

$$\Delta T = \frac{T_{95, \text{ablation}} - T_{95, \text{baseline}}}{T_{95, \text{baseline}}} \quad (3)$$

$T_{95}$  is the time it takes for 95% of people to complete the evacuation. We use  $T_{95}$  instead of the average to avoid individual extremely slow cases from skewing the mean. Other indicators include: casualty difference  $\Delta N$ , which is the difference between the number of people whose FED reaches the incapacitation threshold and the baseline; safety margin distribution, which is the distance between an individual's FED value and 1 when they reach the exit—smaller margins mean they're closer to incapacitation; and exit usage balance, measured by the coefficient of variation of people passing through each exit, reflecting how uneven exit use is due to incomplete information. These four indicators cover three levels:  $\Delta T$  measures overall efficiency deviation,  $\Delta N$  and safety margin distribution measure the severity of consequences, and exit balance measures how the deviation is distributed spatially. Only by looking at all three can we distinguish between "slower evacuation" and "evacuation becoming more dangerous," which are two different types of prediction biases.

To quantify a channel's contribution, the channel's marginal contribution is defined as the proportion of the deviation caused by downgrading that channel alone relative to the deviation caused when all three channels are downgraded together; the channel interaction term is defined as the difference between the combined downgrade deviation and the sum of the individual channel deviations, and a positive interaction term indicates that there is a synergistic amplification between channels.

## 4. Simulation Platform and Ablation Experiment Design

### 4.1 Simulation Core Coupling Fire, Smoke, And Occupants

The experimental setup in this paper is a multi-agent evacuation simulation platform that couples fire, smoke, and human factors. The platform discretizes the building into a cellular space with a 0.5 m resolution. The flame field evolves using a non-linear growth CA model, and it is calibrated by cross-checking with typical fire scenarios simulated with FDS. The smoke field simulates the diffusion and settling of smoke, providing input for visibility attenuation. Each agent has a state machine with four states: normal walking, obstacle avoidance, detouring, and incapacitation. Path decisions use a risk-weighted approach, combining target distance, fire risk, and smoke risk weighted linearly into a movement cost. Agents move in the direction of the lowest cost while updating their environmental observations. Damage accumulation for agents exposed to fire is calculated using the FED model, with FED reaching 1 indicating incapacitation. The FED value at the exit time represents the individual's safety margin [7]. The method of quantifying local fire risk along the path is consistent with Cao et al[12], and overall, it follows a multi-agent modeling paradigm [23].

We chose a discrete cellular space instead of a continuous space model because ablation experiments need to be repeated dozens of times across 198 combinations formed by 18 scenarios and 11 operating conditions, and the millisecond-level single-step cost of the CA model makes this scale computationally feasible. We opted for continuous FED rather than a binary casualty flag because comparing the distribution of safety margins requires a continuous measure, while binary flags would hide individual differences that are high-risk but not fatal.


### 4.2 Logic of the Ablation Experiments

The experiment was carried out in three stages. Stage one established an all-knowing baseline to serve as a common reference for all ablation conditions. Stage two involved single-channel ablation, where perception, cognition, and delay were each reduced from the all-knowing level to partial and scarce levels, giving 3 channels  $\times$  2 levels = 6 conditions, isolating the independent contribution of each channel. Stage three involved combined ablation, where under the scarce level, channels were paired, or all three were downgraded simultaneously, giving 4 conditions to capture channel interactions. All scenarios are paired under the same fire source, number of people, and exit configurations, using the same random seed sequence, so the differences between scenarios only come from the information channel settings. The overall process of the ablation experiments is shown in Figure 1. The 'partial' level corresponds to the basic information ability of a real crowd, and the 'lacking' level corresponds to a complete absence of information ability. These two levels are used as a contrast to capture the non-linear bias as information degrades; the combined ablation is carried out at the 'lacking' level to check if there are super-additive interactions between channels.

### 4.3 Benchmark Scenario and Experimental Matrix

The baseline scenario is a single-story rectangular space measuring 24 m  $\times$  15 m  $\times$  3 m, with a grid resolution of 0.5 m. There are three candidate fire source locations in the space (near the exit side, the center of the space, and the far exit side) and two exit configurations (2 exits, 4 exits). The crowd size has three levels (60, 120, 240 people), covering low, medium, and high density. The fire source location, exit configuration, and crowd size make up 18 basic scenarios, and each basic scenario is combined with 11 ablation conditions. The values for the three-channel parameters are shown in Table 1, and the scenario and ablation condition matrix are shown in Table 2. A three-factor orthogonal combination can separately isolate how the relative position of the fire source, exit redundancy, and crowd density affect the bias under incomplete information. The scene size is based on the scale of a typical single-story commercial space. Each basic scenario is combined with 11

operating conditions (1 baseline, 6 single-channel, 4 combinations), resulting in a total of  $18 \times 11 \times 30 = 5940$  paired runs.



**Fig. 1.** Decomposition of the omniscience assumption into three channels and the ablation experiment procedure

**Table 1**  
Definitions of parameter levels for the three channels

| Channel    | Parameter                          | Omniscient                      | Partial                                             | Deprived                                       |
|------------|------------------------------------|---------------------------------|-----------------------------------------------------|------------------------------------------------|
| Perception | Visibility radius $r_v$            | Unrestricted global perception  | Decreases with smoke concentration; lower limit 2 m | One cell; wall-following search                |
| Cognition  | Known exits $E_{\text{known}}$     | All exits and fire source known | Familiar subset of exits; fire source unknown       | Empty set; random movement or following others |
| Delay      | Information update interval $\tau$ | Updated at every step           | Fixed at 10 s                                       | No updates throughout evacuation               |

**Table 2**  
Scenario and ablation matrix

| Dimension            | Values                                                                                                                                                                                                                                         |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fire-source location | Near exits/center of space / far from exits                                                                                                                                                                                                    |
| Exit configuration   | 2 exits / 4 exits                                                                                                                                                                                                                              |
| Population size      | 60 / 120 / 240 occupants                                                                                                                                                                                                                       |
| Ablation condition   | Baseline; partial perception; deprived perception; partial cognition; deprived cognition; partial delay; deprived delay; deprived perception + cognition; deprived perception + delay; deprived cognition + delay; all three channels deprived |

#### 4.4 Statistical Reporting and Convergence

Each scenario was run 30 times in pairs using the same random seed sequence, recording the T95, number of casualties, and exit usage for each run, and reporting the results as the mean and

95% confidence interval. The number of runs was determined based on Ronchi and others' discussion on the number of runs for stochastic evacuation models: we first checked the stability of the T95 mean under the baseline scenario, and if the mean drifted by more than 2% between consecutive batches, we added more runs until it converged [19]. Comparisons between operating conditions are based on whether the confidence intervals overlap, only reporting the interval relationships without claiming significance tests. The paired design uses the same random seed sequence across all conditions, with identical individual initial positions, random disturbances, and fire evolution disturbances, so that differences between conditions can be cleanly attributed to the information channel settings.

T95 is defined as the time it takes for 95% of individuals to evacuate. If the evacuation rate in a certain scenario is below 95%, then T95 is set to the maximum simulation time of 450 s as a right-censored value. In this case,  $\Delta T95$  only reflects the proportion of censored data and is not interpreted as a time multiple. The severity of the scenario is determined by the evacuation rate and the number of casualties.

## 5. Results and Analysis

This chapter is organized by research questions: 5.1 Check the omniscient baseline and show the bias magnitude for single-channel ablation (RQ1); 5.2 Answer RQ2 about channel interactions and dominant shifts through combined ablation, spatiotemporal distribution, and dominant channel classification; 5.3 Compare the three types of information states: general, practical, and incomplete. 5.4 Answer RQ3 by transforming the bias decomposition into the basis for prediction adjustments in a fully informed model and the investment priorities for the evacuation information system, and provide a unified mechanism explanation. All results come from 18 scenarios, 11 working conditions, with 30 paired runs per condition (a total of 5,940 runs), reported as means with 95% confidence intervals; when the escape rate for a condition is below 95%, T95 is censored at about 450 s. In such cases, the combination is mainly judged by escape rate, casualties, and average evacuation time of those who escaped, and  $\Delta T95$  only notes the proportion censored.

### 5.1 Omniscient Baseline Check and Single-Channel Ablation

The baseline scenario is used to verify the platform's output under normal conditions without information degradation. In 30 paired runs, for the 2-exit, 120-person, central fire source baseline, T95 was 76.1 s (95% confidence interval 74.7–77.5 s), while for 4 exits of the same size it was 38.1 s (37.2–39.0 s). The evacuation rates for both baseline types were 100%, with zero casualties, which is in the same ballpark as evacuation times in classical CA model literature [5]. Across all 18 scenarios, baseline evacuation rates ranged from 99.7% to 100%, with only the worst-case scenario (2 exits, near fire, 240 people) averaging 0.73 casualties. This sets the premise for all subsequent deviation calculations: if the baseline deviates from literature scales, the ablation deviations lose their reference point. Baseline verification results are shown in Table 3.

**Table 3**

Omniscient baseline check results (n = 30)

| Scenario                                       | Mean T95 / s | 95% confidence interval / s | Order of magnitude in reference literature / s                  |
|------------------------------------------------|--------------|-----------------------------|-----------------------------------------------------------------|
| 2 exits; 120 occupants;<br>central fire source | 76.1         | 74.7–77.5                   | Tens to over one hundred<br>seconds; same order of<br>magnitude |
| 4 exits; 120 occupants;<br>central fire source | 38.1         | 37.2–39.0                   | Tens to over one hundred<br>seconds; same order of<br>magnitude |

Single-channel ablation addresses RQ1. The biases of the six conditions relative to the baseline are summarized in Table 4.

**Table 4**

Relative evacuation-time bias  $\Delta T$  under single-channel ablation (pooled across 18 scenarios;  $n = 540$  pairs)

| Condition           | Mean $\Delta T$ / % | 95% confidence interval of $\Delta T$ / % | Difference in casualty count $\Delta N$ | Change in margin distribution |
|---------------------|---------------------|-------------------------------------------|-----------------------------------------|-------------------------------|
| Partial perception  | -0.10               | -0.28–0.08                                | +0.02                                   | -0.001                        |
| Deprived perception | 2.93                | 0.53–5.33                                 | +0.52                                   | -0.003                        |
| Partial cognition   | 1.28                | 0.73–1.83                                 | -0.02                                   | +0.001                        |
| Deprived cognition  | 28.23               | 24.45–32.01                               | +0.58                                   | -0.014                        |
| Partial delay       | -0.14               | -0.34–0.06                                | +0.01                                   | -0.001                        |
| Deprived delay      | 17.14               | 11.06–23.22                               | +1.55                                   | -0.002                        |

The real experimental results show clear channel layering. Cognitive scarcity is the only independent bias with a significant magnitude across scenarios, with the pooled T95 across 18 scenarios increasing by an average of 28.23% (95% confidence interval 24.45%~32.01%). The cognitive partial quota is only 1.28%, suggesting that whether an export prior exists is far more important than how complete it is. Biases in perception and latency partial quotas do not exceed 0.2%, while scarcity quotas are 2.93% and 17.14%, respectively. The wider confidence interval for latency scarcity is due to right-censoring in the scenario with 2 outputs and 240 people. Breaking it down by group size, we can see that perceived scarcity bias goes from 0.00% at 60 people and 0.09% at 120 people up to 8.69% at 240 people. Delay scarcity rises from 0.03% and 0.09% up to 51.30% (on average, 240 people have 4.65 more casualties, with 17% right-censored in pairs), while cognitive scarcity is the highest across all three density levels (28.12%, 23.06%, 33.52%). This result revises the initial expectation: low-density scenarios are also driven more by cognition than perception, because low-density crowds leave before smoke develops, leaving little opportunity for the perception channel to play a role; the delayed channel doesn't increase just because exit configurations get more complex—it only overtakes when there are fewer exits, higher density, and congestion blocks alternative routes. The mean safety margin changes for six single-channel conditions are all within 0.014, showing that single-channel degradation mainly affects evacuation efficiency rather than individual harm, with its consequences mostly emerging cumulatively in combined scenarios.

## 5.2 Channel Interactions, Spatiotemporal Distribution of Bias, And Changes in the Dominant Channel

The combined ablation addresses the channel interaction part of RQ2. Table 5 shows the deviations of four combined scenarios from the baseline and the interaction terms. Considering that some combined scenarios have escape rates below 95% and T95 is right-censored, Table 5 also reports escape rates, casualties, and their super-additive interaction amounts, and uses the pairwise interaction of average evacuation time per escapee as a robust measure for the time dimension.

The real results depict the structure of synergy amplification more accurately than the original hypothesis: synergy assumes missing cognitive channels. When cognition is intact, even if both perception and delay are lacking (perception-delay lacking), the escape rate still reaches 99.3%, with casualty interactions close to zero and the interaction term for average evacuation time only 0.3%. The two dynamic information channels are almost interchangeable when exit priors are complete, producing no synergy. Once cognition is lacking, combined with a lack of perception, strong super-additivity appears: every single channel on average only causes about 1.2 casualties, but together they cause 63.89 casualties, with an interaction effect of 62.70; the interaction term for average evacuation time of those escaping is 230.9% (95% confidence interval  $\pm 10.3\%$ , not including zero), and the escape rate drops to 58.0%. The synergy from cognitive delay is weaker but in the same

direction (interaction casualties 3.52, escape rate 96.9%). The simultaneous lack of all three channels and the lack of perception and cognition are almost equivalent (escape rate 58.0%, casualties 63.9 people), which suggests that once cognition and perception are already gone, adding delay only gives limited additional impact. It's worth noting that in the table,  $\Delta T95$  reaches a whopping 672% when 90% of pairs are right-censored, mainly due to the 450 s censoring upper limit, and this doesn't reflect the actual multiple of evacuation time. The strength of synergy should be based on the interaction between casualties and average evacuation time; a positive interaction term means the combined downgrade consequences are greater than the sum of downgrades in each channel separately, i.e., the bias isn't simply additive, and single-channel experiments will underestimate the true deviation in scenarios where information is incomplete.

**Table 5**


Combined ablation and channel interactions (pooled across 18 scenarios; escape rates, casualties, and interactions in mean evacuation time provide the primary evidence for right-censored conditions)

| Condition                       | $\Delta T95$ / %<br>(censored proportion) | Escape rate / % | Casualties/ persons | Sum of single-channel casualties/persons | Casualty interaction/persons | Interaction in mean evacuation time / % |
|---------------------------------|-------------------------------------------|-----------------|---------------------|------------------------------------------|------------------------------|-----------------------------------------|
| Deprived perception + cognition | +672.2 (90% censored)                     | 58.0            | 63.89               | 1.19                                     | +62.70                       | +230.9                                  |
| Deprived perception + delay     | +17.1 (6% censored)                       | 99.3            | 1.60                | 2.16                                     | -0.57                        | +0.3                                    |
| Deprived cognition + delay      | +97.9 (20% censored)                      | 96.9            | 5.74                | 2.22                                     | +3.52                        | -2.9                                    |
| All three channels deprived     | +671.0 (90% censored)                     | 58.0            | 63.90               | 2.78                                     | +61.11                       | +231.5                                  |


Beyond the summary metrics, the spatial and temporal sources of deviations can be pinpointed by the evacuation time of each grid and the spatial distribution of individual FED accumulation, allowing us to tell whether the deviations come from initial route selection errors or from worsening conditions during evacuation. The spatial and temporal distribution of deviations is shown in Figure 2, with darker colors indicating longer cumulative stay times in that grid compared to the fully informed baseline. In the benchmark scenario (Figure 2a, 2 exits·120 people·central fire source), only cognitive deficiencies caused noticeable congestion in the converging corridor to the exits and the detour areas on both sides of the fire, while spatial deviations from perception and delay deficiencies were almost zero. In the worst-case scenario with 2 exits·240 people·fire source near the exit (Figure 2b), deviations appeared across all three channels, with dark spots concentrated in the initial crowd areas far from the exit, the detour areas on both sides of the fire, and the exit queuing area, reflecting the combined effects of initial wrong route choices, limited perception, and queuing.

By further classifying the 18 basic scenarios according to the dominant bias channel, we can address the question of dominant channel shifts posed in RQ2. The rule for determining the dominant channel is: if the bias caused by the lack of a single channel exceeds that of other channels by more than 20%, the scenario is considered dominated by that channel; otherwise, it is judged as a combined effect of multiple channels. The classification results can be seen in Figure 3.

The tested dominant channels show a clear two-level structure. Out of 18 scenarios, 16 are dominated by cognitive channels, covering all low- to medium-density scenarios and all four-exit scenarios. The cognitive deficiency bias is several to dozens of times higher than other channels. The only two delay-dominated scenarios are both high-density setups with 240 people and two exits (under a central fire source, the delay bias is 12.4%, surpassing cognitive at 4.9%; in near-fire scenarios, delay, cognitive, and perception biases reach 295%, 153%, and 50%, respectively).



**Fig. 2.** Spatial distribution of evacuation-time bias by grid cell under deprivation of each of the three channels relative to the omniscient baseline: (a) benchmark scenario (2 exits, 120 occupants, central fire source); (b) most adverse scenario (2 exits, 240 occupants, fire source near exits)



**Fig. 3.** Scenario dependence of the dominant bias channel (18 basic scenarios)

This structure offers two adjustments to the initial expectations: first, low-density scenarios are also mainly guided by cognition rather than perception, because in low-density crowds people have already evacuated before the smoke develops, leaving little chance for the perception channel to

play a role; second, the delay channel doesn't increase just because the exit layout is more complex—it only overtakes when there are fewer exits, higher density, and congestion blocks rerouting spaces. In the four-exit scenario, due to information redundancy and quick evacuation, the delay deviation is always close to zero. The perception channel never dominates any scenario on its own; its role becomes more noticeable in high-fire, high-density situations, but still mainly works through interaction with cognition and delay. The shift of the dominant channel means that different building scenarios need to prioritize strengthening different information systems.

### 5.3 Quantitative Comparison of General, Realistic, and Incomplete-Information Conditions

The ablation mentioned earlier degraded each channel using the omniscient baseline as a reference. In this section, we categorize evacuation scenarios into three information states for quantitative comparison: the general case, which is the omniscient baseline and corresponds to the perfect information assumption in mainstream models; the realistic case, where all three channels are only partially available (visible radius decreases due to smoke, only familiar exits are known, information updates roughly every 10 seconds), corresponding to people with basic information abilities; and the incomplete information case, where all three channels are severely limited, corresponding to the worst scenario with strangers, nearly zero visibility, and no information updates. The three scenarios are strictly paired under the same seed sequence, with results shown in Table 6 and Figure 4.

The results show a two-stage structure. Generally, they're close to the real height: across all 18 scenarios, the actual relative T95 compared to an all-knowing model only deviated by 1.31% (95% confidence interval  $\pm 0.57\%$ ), the escape rate was about 100%, and the average casualties barely rose from 0.04 to 0.07 people. In other words, when basic informational ability is retained, the time error of an all-knowing model doesn't exceed roughly 8%, and it doesn't change the safety conclusion. The jump from reality to incomplete information is more like a cliff than a gradual change: the escape rate drops to 58.0%, average casualties increase to 63.90, and T95 hits the 450 s cap due to a lot of paired right-censoring. The cost of assuming an all-knowing model isn't uniformly slightly optimistic—it's concentrated in scenarios where informational ability is completely missing.


Density affects the direction of adjustments, but the magnitude is within about 5%: at low density, things are slightly slower (60 people: +5.77%), while at high density, some locking at exits reduces door sway, making it slightly faster (240 people: -4.40%). Extra exits provide a buffer for missing information: with the same lack of information, escape rate with 4 exits (75.0%) is higher than with 2 exits (41.0%) and casualties are lower. Overall, the all-knowing model is a good approximation for typical real-world scenarios, and predictions and information investments should focus on worst-case situations with unfamiliar crowds, few exits, high density, or proximity to fire.

**Table 6**

Comparison of general, realistic, and incomplete-information conditions (pooled across fire-source locations and exit configurations within each density level;  $n = 180$  pairs per cell)

| Occupant density              | T95:<br>general/realistic/incomplete / s | Escape rate:<br>general/realistic/incomplete / % | Casualties:<br>general/realistic/incomplete / persons | Correction factor:<br>realistic/general |
|-------------------------------|------------------------------------------|--------------------------------------------------|-------------------------------------------------------|-----------------------------------------|
| Low density: 60 occupants     | 49.5 / 52.3 / 376.8                      | 100 / 100 / 75.7                                 | 0 / 0 / 14.4                                          | 1.057                                   |
| Medium density: 120 occupants | 57.1 / 58.4 / 447.7                      | 100 / 100 / 57.3                                 | 0 / 0 / 47.4                                          | 1.022                                   |
| High density: 240 occupants   | 81.2 / 77.1 / 450.0                      | 99.9 / 99.9 / 41.1                               | 0.13 / 0.21 / 129.9                                   | 0.949                                   |
| All scenarios                 | 62.6 / 62.6 / 424.8                      | 100 / 100 / 58.0                                 | 0.04 / 0.07 / 63.9                                    | 1.000                                   |

Note: Because escape rates under the incomplete-information condition are below 95%, most T95 values are right-censored at 450 s. They are displayed for comparison only and must not be interpreted as multiples of evacuation time. Severity is assessed using escape rates and casualties.



**Fig. 4.** Comparison of T95, escape rate, and casualties among general, realistic, and incomplete-information conditions at three occupant-density levels

#### 5.4 Mechanistic Explanation and Application of Results to RQ3

The expected differences in biases across the three channels can be explained according to the timing of their influence on evacuation decisions, which also corresponds to the different paths through which these biases get amplified. The cognitive channel affects the initial route choice before the evacuation starts: individuals with insufficient cognition tend to choose directions randomly or follow the crowd. In situations where the fire blocks an exit, or the exit layout is asymmetric, a wrong initial choice forces people to backtrack, dramatically increasing the cost of the path, and this cost can't be offset by experience in unfamiliar settings. Perceptual channels act on path correction during evacuation: smoke spread causes the visible radius to continuously shrink, making it impossible for individuals to promptly detect changes in exit status and approaching danger, resulting in delayed correction actions that compound existing cognitive errors. Therefore, the impact of perceived deficiency increases in sync with smoke gas development speed and crowd density, directly driving FED accumulation at high density. Delay channels act on the time rhythm of information refresh: expired snapshots cause individuals to move toward areas that have already deteriorated, with effects depending on environmental change rate and error-correctable space. At low densities, individuals still have room to reroute, and delays are partially absorbed; at high density, congestion alone blocks the possibility of rerouting and the consequences of delays become truly apparent. When all three channels degrade at the same time, initial routing mistakes, the inability to correct them midway, and continuous reaction delays can all link together within the same evacuation segment, so the deviation might end up being greater than the sum of the deviations of each channel. This kind of super-additivity isn't just a statistical illusion: moving continuously in the wrong direction increases both exposure time and backtracking distance, which together drive up the cumulative FED, forming a physical source of synergistic amplification.

For the results of RQ3, there are two application paths for the ablation findings. The first is prediction correction: if a certain type of building scene is confirmed through ablation to have bias dominated by a specific channel, then the evacuation time output by the fully informed model shouldn't be used directly as a safety criterion. Instead, it should be adjusted according to the marginal contribution of that channel, with the correction coefficient corresponding exactly to the bias decomposition structure given in this paper. For typical real-world scenarios, the correction coefficient is around 1, and only in the worst-case situations—unfamiliar layouts, few exits, high density, or proximity to fire—does it need to be significantly increased. Secondly, it's about

information investment priorities: The ablation results from 18 scenarios provide a quantitative basis for prioritization — most conventional scenarios and public buildings mainly visited by strangers are dominated by the cognitive channel, so evacuation signs and exit guidance should be prioritized; places with few exits and high density are guided by the delay channel, so the refresh and broadcast intervals of environmental information should be shortened first; in high-density areas close to fire, the contribution of the perception channel rises significantly, requiring visibility support and smoke control to be strengthened at the same time. The ablation results can also be used to identify the most unfavorable information conditions, which helps in prioritizing emergency drills and facility inspections.

## **6. Conclusions and Future Work**

This article tackles the commonly implicit omniscient assumption in evacuation simulations and carries out a complete study from conceptual formalization to experimental verification: it breaks down the omniscient assumption into three orthogonal information channels—perception, cognition, and delay—describes them uniformly with an information state triplet, and defines a gradual degradation operation from omniscience to scarcity. On a coupled fire-smoke-people platform, we conducted ablation experiments across 18 scenarios and 11 conditions, with 30 paired runs per condition (a total of 5,940 runs), and established a multi-level metric system based on evacuation time deviation, casualties difference, safety margin, and exit balance. The results show that the cognitive channel is the most stable independent source of bias. At low to medium density, cognitive deficiencies increase T95 by about 18% to 31%. The effects of the perception and delay channels grow stronger as crowd density increases and exits decrease, reaching 8.7% and 51.3%, respectively at high density with 240 people. Channel synergy relies on cognitive deficits and shows strong super-additivity—when cognitive deficiencies are combined with perception deficiencies, the average casualties jump from about 1.2 in the two single channels to 63.9, and the interaction effect on average evacuation time reaches 230.9%. In 16 out of 18 scenarios, the cognitive channel dominates, while only 2 high-density, low-exit scenarios are dominated by the delay channel. The dominant channel shifts depending on the scenario conditions. A comparison of the three states—general, actual, and information-deficient—further shows that omniscient models are a pretty good approximation for regular real-world scenarios with basic information capabilities (T95 deviation around 1%, consistent safety conclusions). The downside mainly shows up in extreme cases where information is completely lacking (escape rate drops to 58.0%, average casualties 63.9 people). This framework turns the vague question of 'is the model too optimistic' into a repeatable, traceable experimental procedure, giving a quantitative basis for correcting predictions in omniscient models and prioritizing investments in evacuation information systems.

This paper has four limitations. First, the decay parameter for the perceptual channel's visual radius is based on literature calibration and hasn't been directly calibrated with real human experimental data. Second, the simulation platform does not include physical contact and pushing at the social force level, so casualty assessments in high-density scenarios might be on the conservative side. Third, the ablation study is based on a single rectangular reference scene, so before applying it to multi-room, multi-story buildings, the scene matrix needs to be expanded. Fourth, the delay channel only models information lag at the individual level and does not include system-level information spread like broadcasting and command, which could partially offset individual information deficiencies. Future work will expand the scene matrix to multi-room, multi-story settings, test the robustness of dominant channel transfer patterns, and incorporate system-level information propagation such as broadcasting and command into the delay channel modeling.

## Acknowledgement

This work was supported by the National Natural Science Foundation of China (42377200), the Educational Teaching Reform Research and Practice Project (NCIST funding) (YDJG202523), the Natural Science Foundation of Hebei Province, China (D2025508013); the Open Fund of the Hebei Cangzhou Groundwater and Land Subsidence National Observation and Research Station (CGLOS-2025-05), the Science and Technology Innovation Fund for Postgraduate students subsidized by Fundamental Research Funds for the Central Universities (ZY20250102), and the State Key Laboratory for Tunnel Engineering (TESKL202435)

## Conflicts of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

- [1] Gwynne, S., Galea, E. R., Owen, M., Lawrence, P. J., & Filippidis, L. (1999). A review of the methodologies used in evacuation modelling. *Fire and Materials*, 23(6), 383-388. [https://doi.org/10.1002/\(SICI\)1099-1018\(199911/12\)23:6<383::AID-FAM715>3.0.CO;2-2](https://doi.org/10.1002/(SICI)1099-1018(199911/12)23:6<383::AID-FAM715>3.0.CO;2-2)
- [2] Wang, Y., Kyriakidis, M., & Dang, V. N. (2021). Incorporating human factors in emergency evacuation: An overview of behavioral factors and models. *International Journal of Disaster Risk Reduction*, 60, 102254. <https://doi.org/10.1016/j.ijdrr.2021.102254>
- [3] Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, 407(6803), 487-490. <https://doi.org/10.1038/35035023>
- [4] Blue, V. J., & Adler, J. L. (1999). Emergent fundamental pedestrian flows from cellular automata microsimulation. *Transportation Research Record*, 1644, 29-36. <https://doi.org/10.3141/1644-04>
- [5] Kirchner, A., & Schadschneider, A. (2002). Simulation of evacuation processes using a bionics-inspired cellular automaton model for pedestrian dynamics. *Physica A: statistical mechanics and its applications*, 312(1-2), 260-276. [https://doi.org/10.1016/S0378-4371\(02\)00857-9](https://doi.org/10.1016/S0378-4371(02)00857-9)
- [6] Burstedde, C., Klauck, K., Schadschneider, A., & Zittartz, J. (2001). Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications*, 295(3-4), 507-525. [https://doi.org/10.1016/S0378-4371\(01\)00141-8](https://doi.org/10.1016/S0378-4371(01)00141-8)
- [7] Hartzell, G. E., & Emmons, H. W. (1988). The fractional effective dose model for assessment of toxic hazards in fires. *Journal of Fire Sciences*, 6(5), 356-362. <https://doi.org/10.1177/073490418800600504>
- [8] Zhai, Y., Han, Z. L., Qu, L., & Han, C. (2025). Research on the coupled simulation of fire and pedestrian-vehicle evacuation in underground parking lot. *Chinese Journal of Underground Space and Engineering*, 21(6), 2217-2226. <https://doi.org/10.20174/j.JUSE.2025.06.36>
- [9] Marlow, F., Jacob, J., & Sagaut, P. (2022). A multidisciplinary model coupling lattice-Boltzmann-based CFD and a social force model for the simulation of pollutant dispersion in evacuation situations. *Building and Environment*, 205, 108212. <https://doi.org/10.1016/j.buildenv.2021.108212>
- [10] Mirahadi, F., McCabe, B., & Shahi, A. (2019). IFC-centric performance-based evaluation of building evacuations using fire dynamics simulation and agent-based modeling. *Automation in Construction*, 101, 1-16. <https://doi.org/10.1016/j.autcon.2019.01.007>
- [11] Xie, J. B., Chen, K. C., Kwan, T. H., & Yao, Q. H. (2021). Numerical simulation of the fire emergency evacuation for a metro platform accident. *Simulation*, 97(1), 19-32. <https://doi.org/10.1177/0037549720961433>
- [12] Cao, R. F., Lee, E. W. M., Xie, W., Gao, D. L., Chen, Q., Yuen, A. C. Y., Yeoh, G. H., & Yuen, R.-K.-K. (2023). Development of an agent-based indoor evacuation model for local fire risks analysis. *Journal of Safety Science and Resilience*, 4(1), 75-92. <https://doi.org/10.1016/j.jnlssr.2022.09.006>
- [13] McGrattan, K. B., Baum, H. R., Rehm, R. G., Hamins, A., Forney, G. P., Floyd, J. E., ... & Prasad, K. (2000). Fire dynamics simulator--Technical reference guide. Gaithersburg: National Institute of Standards and Technology, Building and Fire Research Laboratory. <https://doi.org/10.6028/NIST.SP.1018e6>
- [14] Liu, M. T., Zhu, W., Wang, Y. F., & Zheng, J. C. (2021). Modeling and simulation of exit selection behavior in pedestrian evacuation based on information perception and transmission. *Sustainability*, 13, 13194. <https://doi.org/10.3390/su132313194>
- [15] Wang, X., Chraïbi, M., Chen, J., Li, R. Y., & Ma, J. (2022). Modeling boundedly rational route choice in crowd evacuation processes. *Safety Science*, 147, 105590. <https://doi.org/10.1016/j.ssci.2021.105590>

- [16] Yen, H. H., Lin, C. H., & Tsao, H. W. (2022). Novel smoke-aware individual evacuation and congestion-aware group evacuation algorithms in IoT-enabled multi-story multi-exit buildings. *IEEE Access*, 10, 119402-119418. <https://doi.org/10.1109/ACCESS.2022.3221757>
- [17] Ji, Y. P., Wang, W. S., Zheng, M. Y., & Chen, S. (2022). Real time building evacuation modeling with an improved cellular automata method and corresponding IoT system implementation. *Buildings*, 12(6), 718. <https://doi.org/10.3390/buildings12060718>
- [18] Mendoza, M. G., Waldburger, L., Lee, J., & Sastry, S. (2026). Hierarchical Generative Agents for Simulating Sequential Human Behavior. *arXiv preprint arXiv:2606.14989*. <https://doi.org/10.48550/arXiv.2606.14989>
- [19] Ronchi, E., Reneke, P. A., & Peacock, R. D. (2014). A method for the analysis of behavioural uncertainty in evacuation modelling. *Fire Technology*, 50(6), 1545-1571. <https://doi.org/10.1007/s10694-013-0352-7>
- [20] Jin, T. (1972). Visibility through fire smoke (III). *Bulletin of Japan Association for Fire Science and Engineering*, 22(1\_2), 11-15. <https://doi.org/10.11196/kasai.22.11>
- [21] Isobe, M., Helbing, D., & Nagatani, T. (2004). Experiment, theory, and simulation of the evacuation of a room without visibility. *Physical Review E*, 69(6), 066132. <https://doi.org/10.1103/PhysRevE.69.066132>
- [22] Zhiming, F., Xingpeng, X., Lixue, J., Xiaolian, L., & Nan, H. (2021). Study on the exit-selecting behavior in underground indoor space with fire using a virtual experiment. *Tunnelling and Underground Space Technology*, 112, 103936. <https://doi.org/10.1016/j.tust.2021.103936>
- [23] Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(Suppl. 3), 7280-7287. <https://doi.org/10.1073/pnas.082080899>
- [24] Proulx, G. (1995). Evacuation time and movement in apartment buildings. *Fire Safety Journal*, 24(3), 229-246. [https://doi.org/10.1016/0379-7112\(95\)00023-M](https://doi.org/10.1016/0379-7112(95)00023-M)