An Agent-based Framework for Evaluating Location-based Risk Score in Indoor Emergency Evacuation

Tianlun Cai[a], Hong Zheng[a], Jiamou Liu[b], Yupan Wang[b] and Vicente Gonzalez[c]

[a]School of Computer Science and Technology, Beijing Institute of Technology, China
[b]School of Computer Science, University of Auckland, New Zealand
[c]Civil and Environmental Engineering, University of Auckland, New Zealand

E-mail: kojiharufish@163.com, jiamou.liu@auckland.ac.nz, hongzheng@bit.edu.cn, ywan980@aucklanduni.ac.nz, v.gonzalez@auckland.ac.nz

Abstract -

The evaluation of indoor risks is a paramount issue in building design and construction. Conventional methods that rely on handcrafted rules or drills are insufficient for this task as they either fail to accurately depict the sophisticated spatial attributes and people's cognition abilities or are not suitable during the design phase of the building. This paper puts forward a novel computational framework with a reinforcement learning-based paradigm to automatically assess the evacuation risk posed by the indoor space through intelligent agents. Our model focuses on the agent's exploration behaviours as it gains knowledge to locate the optimal path from an initial location to an exit. The cost of this knowledge acquisition process is then used to capture the risk posed by the initial location of the building. The work aims to shed new light on utilizing agent-based techniques in evaluating building safety.

Keywords -

Indoor risk evaluation; Reinforcement learning; Agent-based evacuation; Computational framework

1 Introduction

The evaluation of indoor risks posed by emergency situations such as fire and earthquake has been a critical issue in building design and construction management. Imagine a situation where a designer is planning the interior space of a building. It is important to identify potential emergency situations and assess the level of risks of the designed indoor layout before commencing development. One of the key safety criteria is how effective the interior layout enables building occupants to evacuate during an emergency. A well-designed interior layout should in principle enable people in the building to quickly find optimal evacuation pathways and escape, thereby greatly reducing human casualty. Yet, to characterise the indoor environment and hazards, a huge amount of variables need to be taken into consideration that includes not only the spatial features of each interior location, placement of structural and non-structural elements, but also people's behavioural and cognitive traits. This demands an objective, cost-effective, and general method to assess indoor risks which differentiate risks of different indoor locations, truthfully reflect the spatial features, and are not biased by human input.

So far, however, most commonly used methods to evaluate indoor space risks rely on fixed sets of well-defined criteria, e.g., the equations in SFPE handbook of fire protection engineering [1]. It is easy to see that limitation with this common practice: In a building with complex indoor layout structures, applying a set of prescriptive rules to evaluate risks is deemed to be too coarse to account for the differences among the spatial features of the locations, and the erratic and complex behaviours of building occupants [2]. Another common way to evaluate evacuation effectiveness is by conducting evacuation drills. However, this approach introduced various issues related to ethical, practical, and financial constraints [3]. For example, drills will hardly recreate the sense of urgency in a real-world emergency and the effectiveness of evacuation drills greatly relies on whether people actively engage in the training. Moreover, conducting evacuation drills has one obvious shortcoming: we can only perform evacuation drills after the entire interior layout has been developed. Thus, it is not a suitable way to reveal potential risks during the design phase [4].

A desirable scheme for assessing indoor evacuation effectiveness during the building design phase must satisfy the following:

1. Firstly, as the evaluation should take place during the design phase where no physical involvement of human participants is possible, the scheme should exploit computerised modeling and simulation techniques to achieve its goals.

2. Secondly, the evacuation process's effectiveness relies on the evacuees' physical and mental ability. Therefore the assessment must be carried out based on the behavioural patterns of evacuees.

3. At the design phase, however, no input will be avail-

able on the specific features of the building occupants. It is therefore desirable for the risk assessment to be independent from the features of any specific evacuee. Similarly, the specific emergency situation that triggers the evacuation may greatly impact the evacuation process. Therefore, the assessment should also be independent from any specific emergency situation. In summary, the output of the assessment would be a form of *risk score* that measures the level of safety guaranteed by the indoor environment given the interior building layout in a generic setting.

4. Then, as an indoor environment has complex layout structure and locations may have vastly different spatial features, it is important to have a location-based risk assessment where risk scores are associated with individual location points. This has the result of a type of "heat map" that visually illustrates areas of potential hazards in the building where building occupants may face a greater level of risks.

5. Lastly, the scheme should be sufficiently general so that it can be applied to different layout plans and reveal their differences.

This paper proposes a novel framework for computing risk scores of indoor locations given a layout plan. The risk score reflects in a generic sense how much the interior layout supports effective evacuation of a building occupant from a specific location spot. The computation is agent-based in the sense that reinforcement learning (RL), commonly seen in the development of artificial intelligence agents, as an integral part of risk assessment. RL has demonstrated wide applicability in improving the performance of AI systems in many fields [5, 6]. Yet, to the authors' knowledge, our work is the first to incorporate RL in building risk assessments. The aptness of the RL paradigm in our setting lies in the fact that we model an evacuee as a reward-driven decision-maker, i.e., an agent who can assess the physical space while finding an optimal evacuation pathway.

In a nutshell, our framework defines the risk score using the efficiency of an agent – which is initialised with no knowledge regarding the building layout – in finding the optimal evacuation path through exploring the indoor environment. This framework has the following advantages:

Firstly, by applying machine learning algorithms on evacuees who have no prior knowledge regarding the indoor layout, the behaviours of the evacuees are generated in run-time through their interactions with the indoor environment. This avoids handcrafted behaviours of evacuees, thereby giving an unbiased evaluation of the effectiveness of evacuation.

Then, as we let the same agent start its exploration from all location points in the indoor environment, the obtained scores can be compared with each other in an objective way. Through this, one can easily generate a unified heat map of the indoor layout.

Thirdly, the proposed risk score takes into account the cognitive aspect of the agent by focusing on the dimension of *knowledge acquisition*. In other words, the risk score of a location point is defined as the cost for an evacuee who starts from the location points to identify the optimal evacuation path through repeated simulated exploratory evacuation. A location that facilitates higher safety is seen as one that incurs the lower cost of knowledge acquisition, while a more risky location is one where an evacuee needs to spend a lot more effort finding the optimal evacuation path. In comparison with other classical methods for evacuating location-based evacuation effectiveness such as computing the distance between a location and the nearest exit, our formulation is more realistic as it takes into account the behavioural aspect of evacuees and thus is a more reliable and robust safety index.

It is important to note that, while our formulation of the risk score is based on a simulation of evacuation behaviours of an evacuee, the aim of this model is not to mimic the actual evacuation during an emergency scenario [7]. To do that, many factors such as dynamics of the physical space during an emergency (e.g., fire), people's reaction to crowd movements (e.g., herding, stampede, etc.), as well as behavioural traits such as panic and altruism need to be addressed. Our work does not address these factors, as they are specific to particular emergency situations.

2 Related work

2.1 Building Safety Evaluation

The most standard practice to evaluate building safety is to adhere to a set of rigid rules which lays out best practices in building design. One of the main weaknesses of such approaches is that it is not able to accurately and deeply depict the erratic and complex behaviour and movement pattern of evacuees in a building under urgent conditions. Therefore, modeling and simulating emergency evacuation is essential to provide valuable insights about building safety and evacuation management [8].

In 1998, Ming [9] proposes a fire safety assurance approach including the fire safety assessment method for high-rise buildings in Hong Kong. In 1999, Ming [10] proposes a fuzzy fire safety assessment approach based on fire risk ranking techniques. The research at this stage carries out an emergency safety assessment on the completed buildings, the purpose is to evaluate the emergency safety risks of the buildings and formulate appropriate rectification plans to ensure the safety of the buildings.

Since the 2000s, researchers have shifted their attention towards the evaluation of building safety in real emergency situations. In 2001, Ellingwood [11] studied the emergency safety of buildings under earthquake disasters. In 2008, Anagnostopoulos et al. [12] studied the post-earthquake emergency safety assessment of the building and provided support for the post-disaster rescue work plan. Carrying out research in a real emergency play an important role in reducing casualties and property losses under various disaster conditions.

In recent years, it has become a new research direction to carry out safety evaluation on design and construction stage. In 2010, Gangolells et al. [13] discussed the safety considerations of building construction at the design stage using a risk analysis method, aiming to reduce construction risks in advance. In 2011, Oien K et al. [14] conducted a research on the theoretical basis of building safety evaluation, which provided an important reference for subsequent research. In addition, based on the behaviour of people in the building under emergency evacuation conditions, carrying out building emergency safety evaluations to serve the safe evacuation of people under emergency evacuation conditions has become the focus of relevant research. In 2018, Bahr [15] conducted extensive discussions on the safety engineering and risk assessment of system based on practical methods in his work.

2.2 Evacuation modeling

We now review theories and models for evacuation simulation.

Flow-based models. Flow-based models use the density of nodes in flows to simulate the features of the people flow. Henderson [16] firstly argued that the behaviour patterns of pedestrian crowds are similar to gases or fluids. Bradley [17] applied Navier–Stokes equations to describe the motion of high-density pedestrian crowds. Helbing et al. [18] summarised that for medium and high-density pedestrian crowds, its motion patterns are very similar to fluids. For instance, people's footprints in snow look similar to streamlines of fluids. Flow-based models could apply the complete network model to develop the optimal evacuation plan in terms of the minimum evacuation time. However, the main restriction of flow-based models is it makes wrong assumptions of people's homogeneity. These assumptions make this type of model difficult to simulate people's different physical abilities, behaviour patterns, and characters. That is, the sociological factors of group decision-making processes that play a crucial rule in all emergency evacuation scenarios could not be simulated and defined in flow-based models.

Cellular automata. Cellular automata evacuation models involve discretization space and model people's evacuation process by single individual cells. One of the earliest cellular automata evacuation models was proposed by Perez et al. [19]. Daoliang et al. [20] applied a two-dimensional cellular automata model to simulate the evacuation dynamics of occupants. Yu and Song [21] proposed a model to simulate pedestrian counter flow in an corridor considering the surrounding environment. Kirchner et al. [22] proposed a stochastic cellular automata model to simulate the friction effects and clogging phenomenons in the crowd during the evacuation process.

Due to the simple shape of grids and predefined behaviour rules of evacuees, cellular automata models hardly to simulate unique characteristics and behaviour patterns for different types of evacuees, like women or kids. Therefore, most of the complex sociological factors among evacuees cannot be simulated in cellular automata models.

Agent-based models. An agent-based model is a system that comprises many intelligent agents which can be used to build heterogeneous characteristics and behaviour patterns. Instead of a global goal, each of the agents has a local goal to achieve [23]. Camillen et al. [24] evaluated different evacuation strategies in closed spatial environments, they demonstrated that due to the dynamic environment, traditional simulation models are difficult to simulate the evacuation process accurately. Only agent-based models are able to capture the dynamic characteristics of certain closed spatial environments properly. More recently, Sumam et al. [25] focused on the impacts of various evacuation behaviours and determined their efficiency in terms of the final evacuation rate. Yang et al. [26] present a prototype that uses agent-based modeling to simulate deep foundation pit evacuation in the presence of a collapse disaster. Şahin et al. [27] proposed an approach which combines a multi-agent model with fuzzy logic to simulate common human and group behaviour during safety evacuation. Kasereka et al. [28] proposed an intelligent Agent-Based Model enabling the modeling and simulation of the evacuation of people from a building on fire. Gonzalez et al. [29] propose a simulation model to find out optimum evacuation routes during a tsunami using Ant Colony Optimization (ACO) algorithms.

3 A New Agent-based Framework for Assessing Risk Score

We describe our computational framework to assess the risk score given input layouts. A prototype of our framework is implemented using the NetLogo platform. NetLogo is a programmable agent-based modeling environment designed and authored by Wilensky [30]. The physical space in a NetLogo environment is realised by a set of *patches* that represent location points. An agent is a special entity in the framework which can be in a patch at any given time instance. The NetLogo platform has built in function that enable us to define the states, perception and decision making functionalities of the agent, before letting the agent to start its simulated runs, i.e., evacuation in the building.

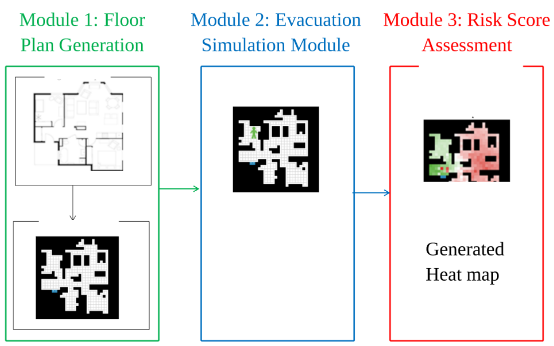

Figure 1. Framework of our approach, with three main modules.

Broadly speaking, our framework consists of three main modules (as shown in Figure1): the floor plan generation module, evacuation simulation module, and risk score assessment module.

First, the indoor space is represented by a digital model that is to be processed by the NetLogo platform. This digital model captures the interior layout and regions that are reachable by the evacuee. The many spatial attributes of the interior layout are represented in the floor plans, such as pathways, exits, corridors, furniture, etc. This digital representation is going to serve as the input to the next module which performs RL algorithm to assess the evacuation risk and generate risk scores.

The second module is the core part of the framework and it simulates the exploratory behaviour of an evacuee in the building. The main idea is we focus on the exploration behaviours of the agent as it gains knowledge in order to locate the optimal path from an initial location to an exit. This corresponds to a process of knowledge acquisition. As discussed above, the cost of the knowledge acquisition process is then used to capture the risk posed by the initial location of the building.

The third module is used to generate risk scores. In particular, we apply the evacuation simulation assuming the evacuee starts from every location points of the physical space. In this way, this module will generate a risk score for every location point. Using these risk scores, one may derive an overall risk score for the entire floor plan, that is, a quantitative measurement of the input floor plans' safety level in terms of emergency evacuation. Moreover, we will generate a heat map based on each patches' risk score.

3.1 Module 1: Floor Plan Generation

The first module performs preliminary processing of floor-plans: we label the inaccessible areas (like walls or large cabinets) and the exits, and then import the modified floor plan. The system will generate a virtual plan base on the coloured floor-plan. Figure 2 shows the coloured original real-world floor-plan of a hospital and the generated virtual plan in Netlogo separately. The black patches represent the inaccessible area, and the white patches represent the accessible area, that is agents could arbitrarily move in white areas. Finally, the blue patches represent the exits.

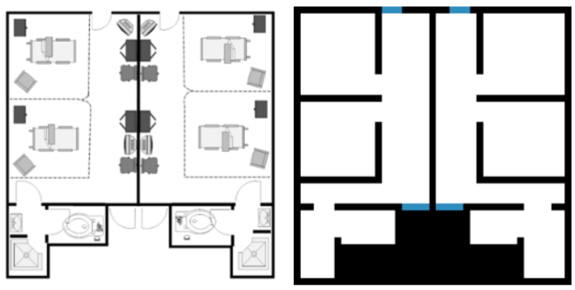

Figure 2. The coloured original real-world floor-plan of a hospital (left) and the digital representation in Netlogo (right).

3.2 Module 2: Evacuation Simulation Module

The main task performed by this module is to assess the evacuation risk posed by the indoor space automatically. Here we use one of the well-known machine-learning paradigm, reinforcement learning (RL), to model a person's cognitive behaviours. An RL agent is a reward-driven decision-maker who's able to adjust behaviours and improve performance based on the external environment through repeated trial-and-error. In this way, RL aims to mimic the cognitive process of operant conditioning in humans and animals. Imagine someone starting from an arbitrary location in a building that aims to escape the building by finding the nearest exit. Suppose that this person has no knowledge regarding the physical space. The person may explore the space by conducting *trials*, i.e., walking in the indoor area until reaching an exit point, or coming to a location that has already been visited in the same trial. Within each trial, certain knowledge is developed by the person that indicates how easy it is to find the nearest exit from this location.

To realise the exploration and learning process above, we formalise the evacuation situation using finite *Markov decision processes* (MDP): An MDP is a tuple $\langle \mathcal{S}, \mathcal{A}, \delta, r \rangle$ where $\mathcal{S}$ is the finite set contain all the states, $\mathcal{A}$ is the finite set contain all the actions, $\delta : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to [0, 1]$ is the dynamic function, $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is the reward function. The goal of MDPs is to determine a policy $\pi : \mathcal{S} \mapsto A$, a function which maps states to actions, with the maximum expected reward:

$$\arg\max_{\pi} \mathrm{E}\left[\sum_{t=0}^{T-1} r\left(s_t, a_t\right)\right] \tag{1}$$

where $s_{t+1} = \delta\left(s_t, a_t\right), a_t = \pi\left(s_t\right)$ and T is a final time step.

In our setting, the MDP represents the indoor environment of the agent. The set $\mathcal{S}$ of states represents the set of all patches in the digital representation of indoor space. Note that there are only three types of patches in our digital floor-plan layout: black (inaccessible area), white (accessible area), and blue (exits). Agents are only able to walk on the white patches. The action set $\mathcal{A}$ consists of four elements: up, down, left, and right. We assume that these actions will deterministically causes the agent to move from one white patch to another patch in the respective direction. If the target patch is a black wall, the agent will stay put in this time step. An policy directs the agent's movements. To realise the knowledge acquisition process, we assume that the MDP is unknown to the agent, and through a model-free RL algorithm, the agent iteratively improves its knowledge regarding the environment by keeping track of a *valuation function*.

This scenario can be handled by a well-established RL algorithm, *Q-learning*. Q-learning has been the most widely-used approach with demonstrated efficiency and reliability guarantee [31]. The key idea of Q-learning is incrementally approximating the valuation (Q) function of each state-action pair based on the rewards received. To be more specific, in each round, the *Q-value* will be updated from Q_t to Q_{t+1} base on old value (i.e., Q_t) and the maximum Q-values of the next state (i.e., s_{t+1}) using a temporal difference mechanism:

$$\begin{aligned} Q_{k+1}\left(s_t, a_t\right) &= Q_k\left(s_t, a_t\right) \\ &+ \alpha\left(r_t + \gamma \max_a Q_k\left(s_{t+1}, a\right) - Q_k\left(s_t, a_t\right)\right) \end{aligned} \tag{2}$$

We now introduce the work flow of the safety assessment module below: A *path* is a contiguous sequence of patches $p_1, p_2, \ldots, p_\ell$ where every patch p_i where $1 \le i < \ell$ is white and p_{i+1} is reachable from p_i by one of the moves. From a patch p, the *shortest path* to exit is the path that starts from p and ends at an exit that contains the least number of patches.

Our agent performs *trials* to explore the indoor area from every white patch p. In each *trial*, the agent starts by initialising the $Q(s,a)$-values to 0 for every patch s and action $a \in \mathcal{A}$. As the agent traverses through the patches, it will update its Q-value as defined above. When the agent arrives at the exit patch, the module checks whether the agent has found a shortest path to the closest exists in terms of its Q-values for each patch-action pair. If no shortest path is identified based on the Q-value, we will start a new trial from the starting patch; on the other hand, if the shortest path has been identified, the module will record the total number of patches this agent has walked through during the all the trials which started from the same initial patch p, this number is denoted by C_p.

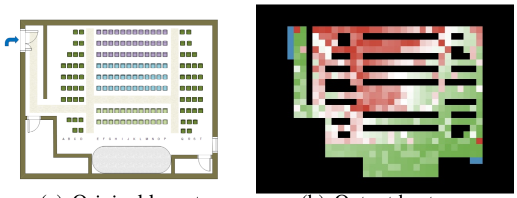

(a) Original layout (b) Output heat map

Figure 3. Original layout and heat map of an theater. Retrieved May 22, 2019, from `https://www.cadpro.com/draw-floor-plans/`

3.3 Module 3: Risk Score Assessment

For any white patch $p \in \mathcal{S}$, let d_p be the length of the shortest path that starts from p. To account for the impact made by different sizes of the floor plan (i.e., the number of patches $|\mathcal{S}|$), we divide C_p by d and set:

$$\sigma_p = C_p / p \tag{3}$$

We define this value σ_p as the *risk score* of the patch p.

We will generate a heat map base on each patches' risk score. We calculate the overall average of σ_p and denote the result σ as the *risk score* of the entire indoor space. We try to use this risk score σ to make the indoor location-based risk quantitative and allow people to judge the configuration of the floor plan as a whole.

4 Experimental Validation

Parameters setup. We apply frequently-used parameters in Q-learning here. In the following experiments, we set ϵ as 20%, α equal to 0.2, γ equal to 0.9, respectively. Here ϵ = 20% means random actions are taking 20% in all actions for the agent's behaviour. We set the path reward -1 to correspond to the time cost during the exploration, and we set the exit-reward to 10 as the reward of finding the exit. Since the Q-learning process is non-deterministic, the choice of the agent's move is not the same in each episode, so we run our model 20 independent episodes for each experiment.

4.1 Case Study 1: A Theater

Figure 3 is a floor plan of a theater[1]. Here, we execute our program on this floor plan and output the heat map. The area near the exit door have a relatively low output risk score (green area), which means such areas are easy to evacuate; meanwhile, the areas far from the exit get high output risk score. This result is matching our expectation and common sense. Moreover, the top left corner is covered with dark red. The reason is there are too many obstacles (i.e., chairs.) around. So it is difficult to quickly evacuate from this area in an emergency situation. If a shooter stormed into this theater and starting hurt people, this dark red area might become a dead-end corner. The heat map shows that our prototype program can judge and measure the safety of different sub-areas. From this heat map, we could get some useful design suggestions: we might need to open an exit door at this red area for safety concerns.

One main advantage of our framework is, we could calculate sub-area risk score base on the risk score of each patch. For example, the risk score of the green left-bottom corner is 471; the risk score of the red left-top corner is 1466; the average risk score value of the whole theater is 848.6. This makes our program has more flexibility and capability. We could only focus on the important part of the building which we are interested in.

4.2 Case Study 2: Auckland Hospital

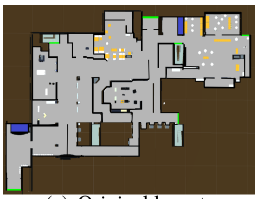
(a) Original layout

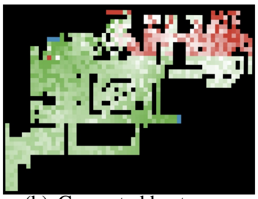
(b) Generated heat map

Figure 4. Original layout and generated heat map of Auckland hospital [32]

Figure 4 shows the original layout and generated heat map of one floor of Auckland hospital, respectively [32]. There are six exits in the original map, as shown in Figure 4(a). Now we remove four of them and generate corresponding heat map, as shown in Figure 4(b), and see what difference it will make. From the graph, we can see the map was separated into two different parts: the left part got two exits and fewer obstacles. Therefore, it has a much lower risk score; the right part has a very complex layout, many obstacles around, some narrow corridors, and dead ends. So we can guess there should be some exits in the red part, just like the original layout. This case study shows that our program does have some ability to understand the structure information embedding in the floor plans.

[1]retrieved May 22, 2019, from `https://www.cadpro.com/draw-floor-plans/`

4.3 Case Study 3: Rational Configuration Design


(a) Original generated floor plan (σ = 652.1)

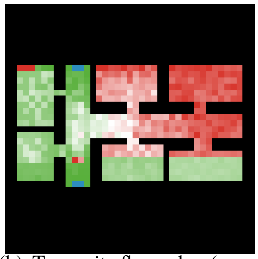
(b) Two exits floor plan (σ = 1655.3)


(c) Narrower corridors floor plan (σ = 803.1)

Figure 5. Rational configuration design

This case study aims to demonstrate that our framework has the potential to provide decision support to designers by allowing flexible adjustment to the indoor layout and observing the resulting risk score. For this, we manually generate an artificial indoor environment that consists of several rooms linked by two corridors. We imagine the situation where a designer is facing a number of design decisions that include setting the number of exists and adjusting the width of the corridors. By applying our model, the designer is able to predict the likely impact to risk scores of the interior space of tuning these parameters.

The number of exits. We first check the difference of the risk score by varying the number of exits. We apply the control variate method here, which means the only thing we change is the number of exits. Figure 5(a) shows the orignal floor plan with six rooms and three exits. The risk score σ is 652.1.

Now we reduce the number of exit to two and rerun our program. We give the output heat map in Figure 5(b). Now the floor plan has a much higher risk score: σ = 1655.3. From the heat maps we generated, we could have some insight view of the whole layout of the current building, like if a certain area has a relevantly high risk score, we might consider adding an exit in this area. The more exit we have, the lower risk score we will achieve. We could set up a proper threshold risk score value to balance the risk score σ and the number of exits.

The width of corridors. We then check how the width of corridors affects the risk score of a floor plan. We reduce the width of the original vertical corridors from four grids to two grids and the original horizontal corridors from four grids to a single grid, as shown in Figure 5(c). By comparing the original floor plan, we find that the narrower corridors floor plan has a higher risk score: σ = 803.1 compared to the original risk score: σ = 652.1. This means narrower corridors make people difficult to evacuate, especially for the two rooms in the middle. This experiment is causing alarm that we must not use the main corridors for storage or rebuild the corridor without permission. We should keep all escape route free and unobstructed evacuation from the premises.

5 Conclusion

The evaluation of indoor risks is a paramount issue in building design and construction. This paper puts forward a novel computational framework to automatically assess the evacuation risk posed by the indoor space through intelligent agents. In particular, we model a person's cognitive process when exploring the indoor space in search for an exit, and then capture the risk using the cost of such process. Such an automated process to evaluate the safety conditions of indoor spaces could help evaluate the risk without deploying safety inspectors or conducting an evacuation drill. Our framework is a cost-effective solution than rule-based risk evaluation or evacuation drills since we perform the evaluation procedure in a simulation environment instead of requiring expensive expertise or conducting an evacuation drill in real buildings. Our framework has also high flexibility because the evaluation procedure could be conduct at any stage of construction, even in the sketch stage, if deemed necessary. Our case studies show that the proposed framework can understand the structure information embedded in the floor plan and offer some reference suggestions for structural design and risk evaluation.

We believe that the proposed method provides a new perspective to evaluating building safety through the lens of computational agents. There are many ways to extend the current work. Future studies could consider a dynamic environment where multiple agents interact through exploring the environment. We also could extend our prototype to the multi-agent system to provide insights into the mechanisms and interactions for panic and crowding under urgent situations. From an application perspective, the idea proposed in the paper can be developed as a plug-in in a building information management (BIM) system that automate the evaluation of evacuation risks.

References

[1] Morgan J Hurley, Daniel T Gottuk, John R Hall Jr, Kazunori Harada, Erica D Kuligowski, Milosh Puchovsky, John M Watts Jr, CHRISTOPHER J WIECZOREK, et al. *SFPE handbook of fire protection engineering*. Springer, 2015.

[2] Margrethe Kobes, Ira Helsloot, Bauke De Vries, and Jos G Post. Building safety and human behaviour in fire: A literature review. *Fire Safety Journal*, 45(1): 1–11, 2010.

[3] Steve Gwynne, Edward R Galea, M Owen, Peter J Lawrence, and L Filippidis. A review of the methodologies used in the computer simulation of evacuation from the built environment. *Building and environment*, 34(6):741–749, 1999.

[4] Xiaoping Zheng, Tingkuan Zhong, and Mengting Liu. Modeling crowd evacuation of a building based on seven methodological approaches. *Building and Environment*, 44(3):437–445, 2009.

[5] Gerald Tesauro, Rajarshi Das, Hoi Chan, Jeffrey Kephart, David Levine, Freeman Rawson, and Charles Lefurgy. Managing power consumption and performance of computing systems using reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1497–1504, 2008.

[6] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.

[7] Gabriel Santos and Benigno E Aguirre. A critical review of emergency evacuation simulation models. 2004.

[8] João E Almeida, Rosaldo JF Rosseti, and António Leça Coelho. Crowd simulation modeling applied to emergency and evacuation simulations using multi-agent systems. *arXiv preprint arXiv:1303.4692*, 2013.

[9] Siu Ming Lo. A building safety inspection system for fire safety issues in existing buildings. *Structural Survey*, 16(4):209–217, 1998.

[10] Siu Ming Lo. A fire safety assessment system for existing buildings. *Fire technology*, 35(2):131–152, 1999.

[11] Bruce R Ellingwood. Earthquake risk assessment of building structures. *Reliability Engineering & System Safety*, 74(3):251–262, 2001.

[12] S Anagnostopoulos and Marina Moretti. Post-earthquake emergency assessment of building damage, safety and usability—part 1: Technical issues. *Soil Dynamics and Earthquake Engineering*, 28(3): 223–232, 2008.

[13] Marta Gangolells, Miquel Casals, Núria Forcada, Xavier Roca, and Alba Fuertes. Mitigating construction safety risks using prevention through design. *Journal of safety research*, 41(2):107–122, 2010.

[14] Knut Oien, Ingrid Bouwer Utne, and Ivonne Andrade Herrera. Building safety indicators: Part 1–theoretical foundation. *Safety science*, 49(2):148–161, 2011.

[15] Nicholas J Bahr. *System safety engineering and risk assessment: a practical approach*. CRC press, 2018.

[16] LF Henderson. The statistics of crowd fluids. *nature*, 229(5284):381, 1971.

[17] GE Bradley. A proposed mathematical model for computer prediction of crowd movements and their associated risks. In *Proceedings of the International Conference on Engineering for Crowd Safety*, pages 303–311, 1993.

[18] Dirk Helbing, Illes J Farkas, Peter Molnar, and Tamás Vicsek. Simulation of pedestrian crowds in normal and evacuation situations. *Pedestrian and evacuation dynamics*, 21(2):21–58, 2002.

[19] Gay Jane Perez, Giovanni Tapang, May Lim, and Caesar Saloma. Streaming, disruptive interference and power-law behavior in the exit dynamics of confined pedestrians. *Physica A: Statistical Mechanics and its Applications*, 312(3-4):609–618, 2002.

[20] Zhao Daoliang, Yang Lizhong, and Li Jian. Exit dynamics of occupant evacuation in an emergency. *Physica A: Statistical Mechanics and its Applications*, 363(2):501–511, 2006.

[21] YF Yu and WG Song. Cellular automaton simulation of pedestrian counter flow considering the surrounding environment. *Physical Review E*, 75(4):046112, 2007.

[22] Ansgar Kirchner, Katsuhiro Nishinari, and Andreas Schadschneider. Friction effects and clogging in a cellular automaton model for pedestrian dynamics. *Physical review E*, 67(5):056122, 2003.

[23] Michael Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2009.

[24] Francesca Camillen, Salvatore Caprì, Cesare Garofalo, Matteo Ignaccolo, Giuseppe Inturri, Alessandro Pluchino, Andrea Rapisarda, and Salvatore Tudisco. Multi agent simulation of pedestrian behavior in closed spatial environments. In *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*, pages 375–380. IEEE, 2009.

[25] Mary Idicula Sumam and K Vani. Agent based evacuation simulation using leader-follower model. 2013.

[26] Weilong Yang, Yue Hu, Cong Hu, and Mei Yang. An agent-based simulation of deep foundation pit emergency evacuation modeling in the presence of collapse disaster. *Symmetry*, 10(11):581, 2018.

[27] Coşkun Şahin, Jon Rokne, and Reda Alhajj. Human behavior modeling for simulating evacuation of buildings during emergencies. *Physica A: Statistical Mechanics and its Applications*, 528:121432, 2019.

[28] Selain Kasereka, Nathanaël Kasoro, Kyandoghere Kyamakya, Emile-Franc Doungmo Goufo, Abiola P Chokki, and Maurice V Yengo. Agent-based modelling and simulation for evacuation of people from a building in case of fire. *Procedia Computer Science*, 130:10–17, 2018.

[29] Eric Forcael, Vicente Gonzalez, Francisco Orozco, Sergio Vargas, Alejandro Pantoja, and Pablo Moscoso. Ant colony optimization model for tsunamis evacuation routes. *Computer-Aided Civil and Infrastructure Engineering*, 29(10):723–737, 2014.

[30] Wilensky. *Wilensky*. Uri, NetLogo (and NetLogo User Manual), Center for Connected Learning and Computer-Based Modeling, Northwestern University, 1999. URL `http://ccl.northwestern.edu/netlogo/`.

[31] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[32] Lin Ni, Vicente Gonzalez, Jiamou Liu, Anass Rahouti, Libo Zhang, and Bun Por Taing. An agent-based approach to simulate post-earthquake indoor crowd evacuation. In *International Conference on Principles and Practice of Multi-Agent Systems*, pages 568–575. Springer, 2018.