The Role of Large Language Models for Decision Support in Fire Safety Planning

Dilan Durmus [1], Alberto Giretti[1], Ori Ashkenazi[2], Alessandro Carbonari[1], Shabtai Isaac[2]

[1] Polytechnic University of Marche, DICEA Department, Ancona, Italy

[2] Department of Civil and Environmental Engineering, Ben-Gurion University of the Negev, Beer-Sheva, Israel

d.durmus@pm.univpm.it, a.giretti@staff.univpm.it, oriashke@post.bgu.ac.il,
alessandro.carbonari@staff.univpm.it,isaacsh@bgu.ac.il

Abstract –

This paper explores the integration of Large Language Models (LLMs), specifically GPT-4, in fire safety planning and knowledge-based systems within the Architecture, Engineering, and Construction industry. Focusing on overcoming challenges in expert systems, it presents an AI-in-the-loop model, illustrating how LLMs enhance decision support. The paper introduces a scenario analysis approach, demonstrating the iterative use of LLMs to enrich expert systems' knowledge bases. Two case studies emphasize the practical application of this approach in fire safety planning, showcasing LLM adaptivity, specialized reasoning, and domain knowledge integration. The study addresses the challenges of LLM-induced hallucinations and emphasizes the need for further research to enhance reliability in technical scenarios. Overall, it contributes to advancing fire safety strategies by leveraging the strengths of LLMs in dynamic building environments.

Keywords –

Adaptive Decision Support; Artificial Intelligence Integration; Fire Safety Planning; Knowledge-Based Systems; Large Language Models (LLMs)

1 Introduction

Fire safety planning is essential in the building industry, with the goal of reducing risks, protecting people, and limiting property damage [1]. As architectural trends evolve and cities expand, there is an increasing need for smart and adaptable fire safety approaches. However, traditional methods face challenges in dealing with the dynamic aspects of real-world building environments [2]. New building codes, materials, and occupancy patterns demand a more flexible fire safety approach that can adapt to the ever-changing landscape. Recognizing the mutable nature of buildings as they adapt to multifaceted influences is fundamental. Underlying factors such as space configurations, material choices and evolving human interactions require not only standard, fixed protocols [1], but also innovative tools that can interactively learn and respond to these changes.

Indeed, Digital twins represent one of the emerging technologies that can deal with such dynamics by means of real-time monitoring, scenario prediction, and proactive measures for optimal functionality. Hence, they cannot rely on predefined, non-adaptive logic-analytic models. [3]. Rather, they must run embedded intelligent tools, capable of processing contextual information and proactively making decisions.

So far, the closest approximation to powerful thinking machines like AI has been expert systems, i.e. knowledge-based systems designed for specific problem-solving tasks. A knowledge-based system (KBS) generally consists of two components: the problem-solving processes (R) and the knowledge base (K). Constructing the reasoning part (R) has been the subject of extensive research and various software tools offer ready-made reasoning engines. Whereas, the almost unexplored challenge lies in creating custom knowledge bases.

Expert systems are inherently knowledge-intensive, representing human expertise in a particular field. Expertise, rooted in extensive domain knowledge, forms the basis of these systems. Constructing each knowledge base is a unique process, starting from scratch and tailored to specific domains or tasks. Acquiring and representing this knowledge is a time-consuming and challenging endeavor that delves deeply into the subject matter, and that must combine a quite high amount of diverse expertise that can hardly be put together on purpose. The underlying philosophy is to create tools that enhance knowledge representation. For this reason, the goal of current research is to develop superior tools for effectively capturing and conveying domain-specific knowledge, acknowledging the importance of a robust foundation for expert systems, which is the point where

AI can provide a valuable contribution to build the knowledge base.

A few studies demonstrate various applications of Artificial Intelligence (AI) in building automation, including expert system design. AI can streamline the design process, improve system operation, and optimize space utilization [4]. Moreover, AI has seen the emergence of powerful tools known as Large Language Models (LLMs), such as OpenAI's GPT-3.5 and GPT-4, that will be investigated in the following of this paper. Trained extensively on text and code, these models understand connections between ideas and offer rich contextual information [5]. This paper explores the potential applications of LLMs in fire safety planning, specifically in "Adaptive Knowledge-Based Systems."

Examining traditional expert systems' performance, the paper highlights the evolving role of LLMs in handling technical complexities and improving fire safety measures in the construction industry. Real-world case studies and National Fire Protection Association (NFPA) guidelines demonstrate how LLMs generate strategic decision-making pathways.

The paper focuses on two NFPA Case Study scenarios, Nightclub Fires (scenario 2 and 3), emphasizing the importance of interacting with AI, using LLMs for strategy generation, and analyzing specific situations.

Highlighting the adaptive nature of artificial intelligence models like ChatGPT-4, this paper underlines their ability to tailor responses to specific conversation contexts. While these AI advancements offer significant potential, caution is advised due to the possibility of hallucinations, where AI models may generate unsubstantiated information, but may even debug the knowledge bases in use and subject to augmentation. The challenges posed by hallucinations and the need for further research to enhance the reliability of AI-assisted fire safety planning will be discussed in the last section of this paper, which aims to contribute to advancing fire safety strategies in the dynamic building industry. Through a comprehensive exploration, it seeks to deepen our understanding of the potential benefits and challenges of adopting AI technologies in fire safety management.

2 Evolution of Decision Support Systems in Technical Domains

Artificial intelligence (AI) systems, including Large Language Models (LLMs) like Generative Pre-trained Transformer (GPT) models, analyze extensive datasets to discern patterns, relationships, draw inferences, provide recommendations, and execute actions. The evolution of conversational AI in the Architecture, Engineering, and Construction (AEC) industry, particularly in Natural Language Processing (NLP) applications, faces challenges with conventional approaches, limiting user interactions [6].

Large Language Models (LLMs), neural networks with extensive parameters, depart from traditional training methods by utilizing self-supervised and semi-supervised learning on vast datasets. Prominent among LLMs are Generative Pre-trained Transformer (GPT) models, notably the OpenAI-developed decoder blocks. GPT models, evolving from GPT-2 to GPT-4 in 2023, operate on transformer-based architectures, capturing statistical patterns in natural language. Renowned for producing coherent, human-like language, GPT models find practical use in diverse applications, including chatbots, content generation, and machine translation, addressing open-ended queries effectively and serving as valuable tools for natural language communication.

The innate development of communication and inference skills in large language models has been confirmed through experiments demonstrating the efficacy of chain-of-thought (COT) prompting. This technique enhances performance across arithmetic, commonsense, and symbolic reasoning tasks, particularly in addressing complex challenges like multi-step math word problems. Language models, when empowered with the ability to generate coherent sequences of intermediate reasoning steps constituting a chain of thought, efficiently solve problems. The impetus behind COT prompting in knowledge management aligns with construction domain objectives. It allows models to deconstruct multi-step problems, allocating additional computational resources to tasks requiring extensive reasoning. Additionally, a chain of thought provides interpretable insights into the model's behavior, elucidating the reasoning process and offering debugging opportunities. Applicable to various language-based tasks, COT reasoning demonstrates general applicability. The elicitation of chain-of-thought reasoning in large language models is achieved effortlessly, emphasizing its limited effectiveness in smaller models. GPT-3's application of chain-of-thought prompting demonstrates favorable comparisons with prior methods, showcasing its performance in addressing intricate problems without fine-tuning on labeled datasets [7].

The construction industry, reliant on diverse information from various stakeholders, faces challenges in integrating, reusing, and efficiently managing information, impacting industry productivity. To address this, effective methods for representing extensive knowledge are crucial [8]. Expert systems, computer programs guided by specific problem-solving knowledge, are pivotal, emphasizing the need for knowledge directing solutions. The term "expert" signifies narrow specialization and substantial competence, benchmarked against human performance [9]. Presently, expert systems serve interactive roles, from smart spreadsheets to financial advisors, relying on knowledge bases elicited through effective techniques [10]. Despite their success, expert systems exhibit less flexibility than humans. Large Language Models (LLMs), like GPT models, enhance expert systems by handling vast unstructured information, resembling human-generated text [11].

Table 1 discusses the potential role of LLMs as a tool for improving knowledge management in the building industry. LLMs expand expert systems by extracting information from diverse, unstructured sources, facilitating knowledge base augmentation and updating for advanced decision-making processes.

Table 1. Using LLMs in combination with expert systems to support decision making

	Expert Systems	Large Language Models
Goal	Supporting decision-making with a specific task	Eliciting additional information to broaden out the expert system
Input	Quantitative and structured data	Any type of (even) unstructured data
Approach	Combination of logical rules and a knowledge base to make inferences	Deriving statistical patterns from evidence and performing reinforcement learning
Method	Symbolic AI	Transformers
Interface	Digital	Answers to questions in human-like language / chain-of-thought
Output	Recommendations and advice for specific scenarios	Arguing scenarios representing the dynamics of complex systems

LLMs may provide a solution to some of the existing difficulties in using expert systems to improve knowledge management in the building industry :

- Data Availability: The large body of documents on which they are trained may provide an answer to the difficulty in obtaining structured data of high quality, which is often not readily available in the construction industry.
- Robustness: The relative resilience of LLMs may allow the drawing of conclusions even when input data is unavailable or inconsistent.
- Adaptivity: Unlike expert systems, generative LLMS have shown an ability to adapt to uncoded situations.

At the same time, some challenges are likely to be faced in the use of LLMs for knowledge management. Primary among those is the fact that such models are prone to hallucination, i.e. they give sound and plausible information that is not true, which reduces system performance and users' expectations. Here, their integration with expert systems may provide a solution, given the ability of such systems to provide consistent inferences based on domain-specific knowledge.

Challenges faced both in the development of expert systems and LLMs are their maintainability, and the costs involved. These stem from the need for large datasets that are complete and accurate and may consequently be difficult and costly to obtain in the first place, and to maintain as changes occur in the construction industry over time.

3 Knowledge Engineering Issue

Knowledge, as employed in knowledge-based systems, encompasses agents' codified experience, serving as the informational foundation for problem-solving. Codification implies the formulation, recording, and preparation of knowledge for practical use. Agents' experience spans well-defined domains, representing fields of interest, and tasks, denoting specific types of work within those domains. Domain knowledge pertains to the general terminology and facts of a domain, while task knowledge involves computational models and facts associated with performing a task.

Two crucial dimensions for describing knowledge are procedural vs. conceptual knowledge and basic, explicit vs. deep, tacit knowledge [12]. Procedural knowledge involves knowing how to perform tasks and encompasses processes, task conditions, task order, required resources, and sub-tasks. Conceptual knowledge relates to knowing that certain relationships and properties exist among concepts, including taxonomies and class membership.

Basic, explicit knowledge is consciously processed and focuses on fundamental tasks, relationships, and concepts. In contrast, deep, tacit knowledge resides in the subconscious and is acquired through experiences, leading to automatic activities. Examples include 'gut feel,' 'hunches,' 'intuition,' 'instinct,' and 'inspiration.' Tacit knowledge is demonstrated in activities like driving a car, where experienced drivers navigate effortlessly based on deep, tacit knowledge, contrasting with learners who lack this intuitive understanding [12].

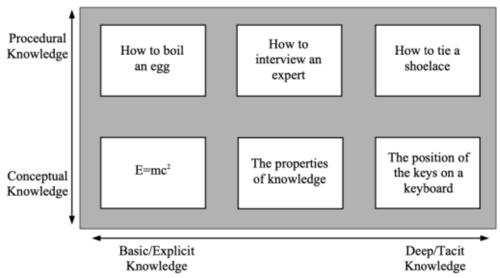

Figure 1.Two important dimensions with which to describe knowledge

3.1 Formulating Expertise

Expertise, typically associated with individuals, goes beyond mere knowledge sources like books and recordings. Experts inject their unique perspectives into the knowledge they possess due to their individual experiences.

3.1.1 Knowledge Acquisition and Implementation Issues

In artificial intelligence and knowledge engineering, knowledge acquisition involves computer systems obtaining the necessary information to perform specific tasks. This process is a three-role participatory design task, involving users, domain experts, and knowledge engineers [9]. The knowledge acquisition process unfolds through four stages at two levels: *knowledge and symbol levels*.

Knowledge-level analysis precedes symbol-level analysis in the knowledge engineering process. During problem identification, unstructured interviews gather crucial sample problems, while conceptualization involves detailed representations of problem-solving processes using protocol analysis. Symbol-level analysis progresses with formalization, translating semi-formal knowledge into a computationally implementable form through defining an abstract language and computational architecture. The structured approach ensures a systematic transition from conceptualization to implementation, considering computational efficiency and user interface aspects. In addition to these stages, knowledge acquisition faces several issues. Experts often struggle to articulate their problem-solving methods, making knowledge acquisition challenging due to tacit knowledge issues. Building common ground between knowledge systems and end-users is complex, and the brittleness of rule-based expert systems outside their limited expertise can lead to nonsensical answers. The expertise scoping issue arises from oversimplified models in knowledge bases, while maintaining consistency becomes crucial for effective validation and reuse of knowledge and problem-solving methods.

3.2 AI in the Loop

In exploring the integration of AI systems into the development of knowledge-based systems based on logical formalism, this article delves into the dual competencies exhibited by current AI systems: domain knowledge and linguistic competence, underpinned by logical-programmatic abilities. Despite the statistical nature of AI systems and occasional performance fluctuations, their expertise presents opportunities to address challenges in constructing knowledge-based systems.

The article proposes an AI-in-the-loop model, extending the participatory design scenario to include AI alongside domain experts and knowledge engineers. This integration operates at both the knowledge and symbolic levels, with a focus on the knowledge level interaction, incorporating psychological aspects of human involvement.

The subsequent sections illustrate examples of AI interaction in the knowledge acquisition process, particularly in the problem identification phase. The focus is on fire emergency management cases, highlighting AI's role in producing areas of expertise and contributing to the initial conceptualization of problems during the analysis of reference situations. This multifaceted integration of AI proves valuable in enhancing the effectiveness and scope of knowledge-based systems in addressing complex real-world challenges.

4 SCENARIO ANALYSIS APPROACH WITH LLMs

This chapter outlines a comprehensive approach to the integration of Large Language Models (LLMs) into the knowledge elicitation process, with a focus on enhancing expert systems (ES) within technical domains. The integration unfolds through two pivotal stages, showcasing the potential of LLMs to contribute contextual knowledge for improved decision support. In the next section 5, the step reported in section 4.1 will be showcased within two different test scenarios; whereas the steps reported in sections numbered from 4.2 to 4.3 have been investigated and discussed, but supposed to be showcased in future research activities.

4.1 Querying LLMs with Embedded Documents

The initial stage involves querying LLMs, particularly Generative Pre-trained Transformer (GPT) models like GPT-4, equipped with embedded documents describing specific scenarios for strategy generation, and analyzing specific situations such as extinguishing a fire or planning an escape route. This process enables contextualization, allowing the LLM to understand scenarios better. By linking documents with user queries in a unified prompt, the LLM generates accurate and context-aware responses. The quality of these embedded documents significantly influences the accuracy of the LLM's responses.

4.2 Enhancing Expert Systems Knowledge Graphs

In the subsequent stage, answers obtained from the LLM as outcomes of the step subject of section 4.1 may be leveraged to enhance the knowledge graph of the expert system. This enhancement can occur directly within the ES's knowledge graph or indirectly through an ontology that underlies the knowledge graph structure. In this context, an ontology provides a structured framework that defines the concepts, relationships, and entities within the knowledge graph, enabling better integration and enrichment of information collected from the language model.

4.2.1 Direct Enhancement with Knowledge Units

Direct enhancement involves incorporating knowledge units gained from querying the LLM directly into the ES's knowledge graph. This meticulous process ensures the integrity of the knowledge graph, with considerations for updating existing nodes and linking newly added knowledge units with relevant scenarios for traceability.

4.2.2 Indirect Enhancement through Ontology

Alternatively, indirect enhancement occurs through an ontology that serves as a conceptual bridge between the LLM and the ES. A semantic ontology defines the ES's knowledge graph and undergoes repeated updates based on information acquired through LLM interaction. The updated ontology, in turn, influences the ES's knowledge graph, fostering a dynamic and evolving system.

4.3 Iterative Processes for Continuous Improvement

By repeatedly applying this integrated process, the ES evolves over time to handle a broader range of scenarios with increased accuracy and relevance. The continuous interaction with LLMs provides a dynamic knowledge base, addressing challenges in information integration and enhancing the adaptability of the expert system.

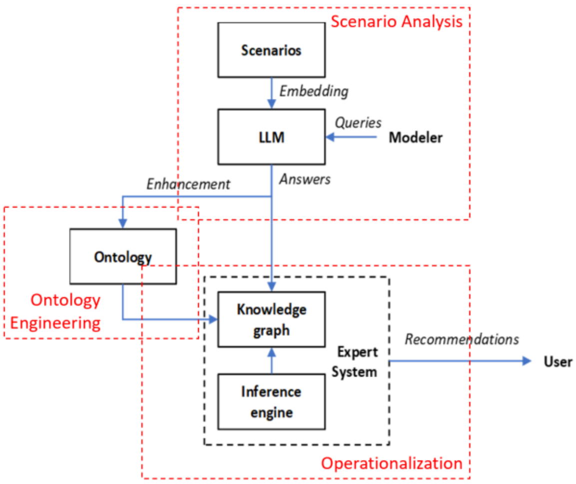

Figure 2. Enhancement of an Expert System using a Large Language Model

At the core of the scenario analysis approach is the interactive engagement with LLMs, pivotal for assessing knowledge and inferences. The next section provides a preliminary evaluation of factors such as consistency, reliability, and adaptivity of information derived from LLMs. A strict scrutiny of metrics of knowledge quality will be set up in later stages of the research, to ensure coherence and logical alignment with established facts, while reliability considerations must deal with contextual alignment. The adaptivity of the expert system is deemed crucial for dynamic adjustments to varying scenarios. A comprehensive evaluation ensures that interactions with LLMs not only enrich knowledge but also to contribute to developing a reliable and adaptive decision support system. As the discussion transitions to practical case studies, the enhanced Expert System (ES), enriched by contextual knowledge from LLMs, forms the groundwork for detailed analyses, exemplifying the practical application of this enriched knowledge in fire safety planning. The subsequent case studies illustrate the synergy between theoretical integration and real-world decision support.

5 CASE STUDY EXAMPLES

In this chapter, two case studies, NFPA Design Fire Scenario 2 and Scenario 3, are presented to emphasize the application of contextual knowledge in fire safety planning in different circumstances. Scenario 2 involves a fire starting in a room lacking a fire extinguishing system but with nearby sand as an available resource, while Scenario 3 focuses on particularly examining fires near exits, requiring a detailed analysis within the NFPA classification system to identify key contributing factors. The importance of engaging in a conversation with AI, using the results to develop the knowledge base of an expert system, and utilizing LLMs for generating strategies and analysing specific situations is demonstrated.

In Scenario 2, a fire starting in a room without a fire extinguishing system is considered, and sand nearby is used as a resource for intervention. The example illustrates a simple yet logical inference process involving generalization and specialization steps (Figure 3).

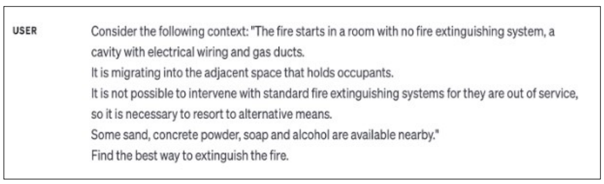
USER

Consider the following context: "The fire starts in a room with no fire extinguishing system, a cavity with electrical wiring and gas ducts.
It is migrating into the adjacent space that holds occupants.
It is not possible to intervene with standard fire extinguishing systems for they are out of service, so it is necessary to resort to alternative means.
Some sand, concrete powder, soap and alcohol are available nearby."
Find the best way to extinguish the fire.

Figure 3. Initial conversation with ChatGPT about the given scenario in NFPA Case Studies

ChatGPT-4 is then asked to provide a fire management strategy, with the assistant's response shown in Figure 4.

ASSISTANT Based on the given context, the best way to extinguish the fire would be to follow these steps:

1. Alert Occupants: Immediately notify the occupants in the adjacent space about the fire and instruct them to evacuate calmly and quickly following the predetermined evacuation routes.

2. Activate Building Alarm: Activate the building alarm system to alert all other occupants on other floors to evacuate as well.

3. Call for Emergency Services: Dial the emergency services number (e.g., 911 or local fire department) to report the fire and provide them with relevant details, such as the location and severity of the fire.

4. Isolate the Area: Close all doors leading to the affected area to contain and restrict the fire's spread.

Adaptivity → 5. Use Sand or Concrete Powder: If available, use sand or concrete powder to smother the fire. Spread a thick layer of sand or concrete powder directly on the base of the fire, covering it completely. This works to cut off the oxygen supply to the fire and help extinguish it.

Hallucination → 6. Soap and Alcohol: Use a mixture of soap and water or alcohol to create a fire retardant solution. Apply the solution onto any combustible or flammable surfaces or objects surrounding the fire, effectively preventing the fire from spreading.

7. Evacuate

Figure 4. Response of ChatGPT4 assistant for the previous user message

ChatGPT-4's notable feature is its adaptivity, showcasing a nuanced understanding of user inputs and tailoring responses to specific conversation contexts. This adaptivity is evident in its ability to suggest fire management strategies, considering available resources and constraints. However, this adaptivity has a flip side, leading to occasional hallucinations where ChatGPT-4 generates information not explicitly in its training data. For instance, in discussing fire retardant foams, an alcohol ambiguity results in a hallucination, as fatty-alcohol stabilizes the foam, but ethanol alcohol is flammable. While ChatGPT-4 can provide helpful responses, users should exercise caution due to the potential inclusion of hallucinatory content. The subsequent analysis delves deeper into domain knowledge regarding soap, water, and alcohol, revealing persistent perceptual anomalies (Figure 5).

USER Explain why soap mixed with water and alcohol creates a retardant solution preventing fire from spreading

ASSISTANT Soap mixed with alcohol can create a fire retardant solution due to the properties of both substances.

Soap contains surfactants, which lower the surface tension of water and help it spread more evenly over surfaces. When mixed with alcohol, it helps the solution become more efficient in coating and wetting the surfaces, making it harder for the fire to spread.

Alcohol, on the other hand, has a high vapor pressure and easily evaporates. When alcohol is mixed with water and applied as a solution, the alcohol evaporates quickly, leaving behind a thin film of water mixed with soap. This film forms a barrier on the surface, preventing oxygen from reaching the fuel source and cutting off the fire's oxygen supply. Additionally, the water in the solution helps to cool the surrounding area, further aiding in fire suppression.

The combination of soap and alcohol creates a fire retardant solution that can effectively coat surfaces, lower the flammability of materials, and prevent the fire from spreading to adjacent areas. It is important to note that a fire retardant solution can only be effective in controlling smaller fires, and it is still crucial to evacuate and call for emergency services.

Figure 5. Response of ChatGPT4 after asking for an explanation of its previous message

Further inquiries for clarification enable a re-evaluation and restoration of contextual accuracy. The challenge posed by hallucination currently hinders direct use of LLMs in technical scenarios, emphasizing the need for additional research and development to mitigate this impact and enhance LLM applicability in technical management contexts (Figure 6).

USER Are you sure that soap mixed with alcohol and water creates a fire retardant solution?

ASSISTANT I apologize for the confusion in my previous response. You are correct that a mixture of soap, water, and alcohol may not create a fire retardant solution.

While soap can help lower the surface tension of water and improve its ability to coat surfaces, the inclusion of alcohol in the mixture can actually increase the flammability of the solution. Alcohol is highly flammable and may act as a fuel source, potentially exacerbating the fire rather than extinguishing it.

In a fire emergency, it is generally recommended to focus on evacuation and alerting emergency services as the primary response. If standard fire extinguishing systems are out of service, utilizing sand or concrete powder to smother the base of the fire can help cut off its oxygen supply and aid in extinguishing it.

Thank you for pointing out the mistake, and I apologize for any confusion it may have caused.

Figure 6. A further request for clarification and response from ChatGPT-4 assistant

Another example explores a second NFPA class no. 2 scenario (scenario 3), addressing fires near exits, particularly in nightclubs. The analysis utilizes the NFPA Case Studies, focusing on factors contributing to such incidents. This examination enriches understanding of fire safety measures and mitigation strategies in vulnerable environments. A user message was sent to ChatGPT-4 (Figure 7).

USER

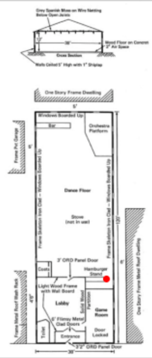

Consider the following context: "A Rhythm Club is a single-story, wood- framed building measuring 120 ft (36.6 m) x 38 ft (11.6 m). The roof and sides of the converted garage building were covered with corrugated metal sheets. The structure contain approximately 24 windows (most of which apparently were shuttered or nailed shut at the time of the fire). The club has only one exit, a door measuring 38 in. (0.9 m) wide. This door opens inward.
Inside the main entrance foyer there is a pair of doors measuring 6 ft (1.8 m) wide, which also open inward. These doors open into a lobby area. Proceeding through the lobby one enter the dance floor, which constituted the majority of the floor space. At the far end of the club, opposite the entrance, there is the orchestra platform and the bar.
The interior of the club is decorated with dried Spanish moss, which is hung on wires from the ceiling joists above the dance floor. The interior walls are constructed of wooden shiplap boards up to 5 ft (1.5 m) from the floor. The floor is composed of wood planks over concrete.
700 patrons and staff and musicians are in the building."

Please plan an escape route if a fire starts near the hamburger stand, which is located adjacent to the only exit from the building and if the fire quickly involves the Spanish moss, spreading rapidly throughout the building, and above the crowd.

Figure 7. User scenario for the analysis of safety planning

The user scenario involves creating an emergency escape plan during a fire at the Rhythm Club. The assistant responds comprehensively, showcasing specialized reasoning, adaptivity, and domain knowledge. The directive, "Given the building's construction, it's essential to guide everyone away from the fire, that is, opposite the direction of the hamburger stand towards the orchestra platform and bar at the end of the club," exemplifies special reasoning. This demonstrates the assistant's nuanced understanding of the building's structural considerations and strategic recommendations based on awareness of physical constraints (Figure 8).

Conversely, the provision, "Since the windows are mostly shuttered or nailed shut, they can be broken open to create additional exits. Critically, keep everyone away from the flames - it's better to break windows furthest from the fire," signifies adaptivity (Figure 9).

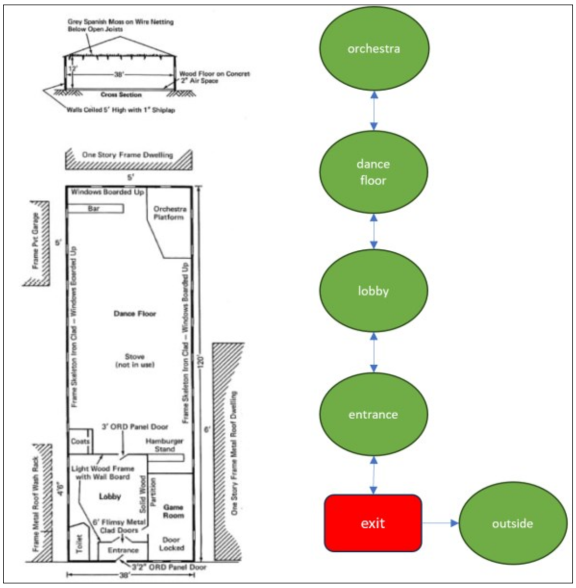

Figure 8. the spatial reference that chatGPT4 was able to produce

Here, the assistant dynamically adjusts the escape strategy based on the environment's specific features. The recommendation to prioritize breaking windows farthest from the fire emphasizes a tactical adjustment for the safety of patrons, staff, and musicians.

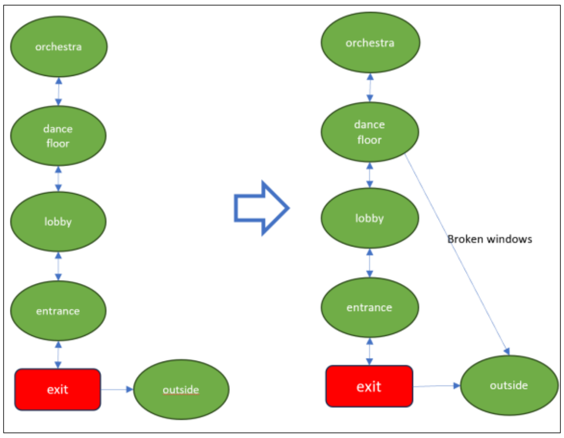

Figure 9. the strategy of creating alternative escape routes by breaking windows

Furthermore, the instruction, "If the fire has spread across the ceiling above the crowd, everyone should slide down on the floor to escape, as the fresher, cooler air will be at ground level," represents domain knowledge. Here, the assistant leverages fundamental insights into fire behaviour and air circulation principles, demonstrating awareness of potential dangers and offering guidance grounded in established knowledge.

When the user seeks clarification on the advice to descend to the floor during a ceiling fire, the assistant provides a scientific rationale. It explains that fire moves upward due to convection, making the air near the ceiling intensely hot and filled with smoke and toxic gases. In contrast, air near the floor remains cooler and cleaner for a longer duration due to cool air's greater density. Sliding down enhances survival chances by ensuring access to more breathable air.

Additionally, the assistant notes that maintaining a lower position improves visibility, as smoke, being lighter, tends to accumulate near the ceiling. Crawling not only keeps individuals below the rising heat but also provides a clearer path to exit, aligning with the well-known safety advice to "Get low and go." Emphasizing the urgency of evacuation, the assistant highlights the rapid accumulation of smoke and heat, even in proximity to the floor. This concise explanation integrates scientific principles and practical considerations, providing a comprehensive understanding of the advocated safety measure (Figure 10).

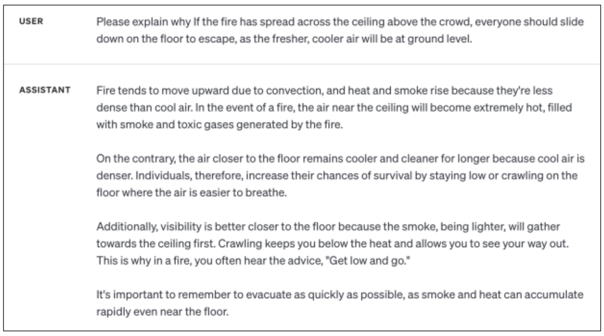
USER Please explain why If the fire has spread across the ceiling above the crowd, everyone should slide down on the floor to escape, as the fresher, cooler air will be at ground level.

ASSISTANT Fire tends to move upward due to convection, and heat and smoke rise because they're less dense than cool air. In the event of a fire, the air near the ceiling will become extremely hot, filled with smoke and toxic gases generated by the fire.

On the contrary, the air closer to the floor remains cooler and cleaner for longer because cool air is denser. Individuals, therefore, increase their chances of survival by staying low or crawling on the floor where the air is easier to breathe.

Additionally, visibility is better closer to the floor because the smoke, being lighter, will gather towards the ceiling first. Crawling keeps you below the heat and allows you to see your way out. This is why in a fire, you often hear the advice, "Get low and go."

It's important to remember to evacuate as quickly as possible, as smoke and heat can accumulate rapidly even near the floor.

Figure 10. Investigating the general knowledge of dominance employed in recommending that occupants lay down on the ground

Finally, inquiries about the accuracy of information regarding cooler and cleaner air near the floor during a fire receive reassurance (Figure 11). The assistant asserts the principle's universal acceptance in fire safety, backed by scientific understanding and common educational practices. While recognizing fire conditions' variability, the assistant underlines the guideline nature of "stay low," not a foolproof safety guarantee. It stresses prompt evacuation adherence to established fire protocols, avoiding hallucination.

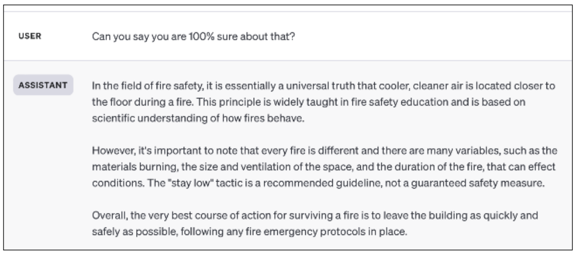
USER Can you say you are 100% sure about that?

ASSISTANT In the field of fire safety, it is essentially a universal truth that cooler, cleaner air is located closer to the floor during a fire. This principle is widely taught in fire safety education and is based on scientific understanding of how fires behave.

However, it's important to note that every fire is different and there are many variables, such as the materials burning, the size and ventilation of the space, and the duration of the fire, that can effect conditions. The "stay low" tactic is a recommended guideline, not a guaranteed safety measure.

Overall, the very best course of action for surviving a fire is to leave the building as quickly and safely as possible, following any fire emergency protocols in place.

Figure 11. Questioning the accuracy of data

Collectively, these components reflect the integration of special reasoning, adaptivity, and domain knowledge within the assistant's response, contributing to the formulation of a comprehensive and informed emergency escape plan in the given scenario.

6 DISCUSSION AND CONCLUSION

The integration of Large Language Models (LLMs), such as ChatGPT-4, presents notable advancements in fire safety planning and knowledge-based systems. These models showcase adaptivity, specialized reasoning, and effective integration of domain knowledge, contributing to enhanced decision support in dynamic scenarios, as evidenced by practical case studies. Despite their potential, challenges, particularly occasional hallucinations, underscore the necessity for cautious utilization and ongoing research to improve reliability. Such feedback can constitute an helpful debug of the current version of the knowledge base, too, by suggesting ways to augment it, e.g. with more accurate definitions of object properties, the existence and nature of relationships that may have been ignored at a previous development stage. In summary, while this paper explores the promising role of LLMs in fire safety, highlighting their adaptability and valuable knowledge contribution, further studies are essential to optimize their application and address challenges like hallucinations. Evaluation metrics for consistency, reliability, and adaptability were initially assessed, and detailed metrics will be provided by the authors in the following study. Exploring the sensitivity of the model to variations in prompt wording and examining different prompts with similar meanings will be crucial in enhancing the robustness of subsequent experiments. Continuous research and development are critical for refining LLMs, ensuring their reliable integration into technical domains, and providing enhanced decision support in fire safety planning. Additionally, future studies are vital for optimizing potential risk mitigation measures for LLMs in knowledge management, including narrowing down their use for specific purposes with embedded documents and incorporating human expertise in the loop to enhance reliability in technical contexts.

References

[1] Kodur V., Kumar P., and Rafi M. M. Fire hazard in buildings: review, assessment and strategies for improving fire safety. *PSU Research Review*, 4(1):1–23, 2019.

[2] Nikolic B. and Dakic J. Knowledge management in the function of risk assessment. *Online Journal of Applied Knowledge Management*, 3(2), 2015.

[3] Grieves M. Intelligent digital twins and the development and management of complex systems. On-line: https://digitaltwin1.org/articles/2-8, Accessed: 22/12/2023.

[4] Panchalingam R. and Chan K. C. A state-of-the-art review on artificial intelligence for Smart Buildings. *Intelligent Buildings International*, 13(4): 203-226, 2021.

[5] Liu Y., Han T., Ma S., Zhang J., Yang Y., Tian J., He H., Li A., He M., Liu Z., Wu Z. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, 1(2):100017, 2023.

[6] Saka A., Taiwo R., Saka N., Salami B.A., Ajayi S., Akande K., Kazemi H. GPT models in construction industry: Opportunities, limitations, and a use case validation. *Developments in the Built Environment*, 17:100300, 2024.

[7] Wei J., Wang X., Schuurmans D., Bosma M., Xia F., Chi E., Le Q.V., Zhou D. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35: 24824–24837, 2022.

[8] Goldstein I. and Papert S. Artificial intelligence, language, and the study of knowledge. *Cognitive Science*, 1(1):84–123, 1977.

[9] Stefik M. Introduction to Knowledge Systems. On-line: https://books.google.it/books?hl=en&lr=&id=SwyjBQAAQBAJ&oi=fnd&pg=PP1&dq=Stefik,+M.+(2014).+Introduction+to+knowledge+systems.+Elsevier.&ots=Qp5vOKRvAU&sig=N5Xwx2bZtlN-DoVy4rWA58MII1w&redir_esc=y#v=onepage&q=Stefik%2C%20M.%20(2014).%20Introduction%20to%20knowledge%20systems.%20Elsevier.&f=false, Accessed: 22/12/2023

[10] Fekri-Ershad S., Tajalizadeh H. and Jafari S. Design and Development of an Expert System to Help Head of University Departments. On-line: https://arxiv.org/abs/1308.0356, Accessed: 22/12/2023.

[11] OpenAI. On-line: https://openai.com, Accessed: 22/12/2023.

[12] Milton N. R. *Knowledge Acquisition in Practice: A Step-by-step Guide*. Springer Science & Business Media, London, 2007.

[13] Lenat D. B. and Feigenbaum E. A. On the thresholds of knowledge. In *Proceedings of the International Workshop on Artificial Intelligence for Industrial Applications*, pages 291–300, Milan, Italy, 1988.