TRANSFORMING BUILDING INDUSTRY KNOWLEDGE MANAGEMENT: A STUDY ON THE ROLE OF LARGE LANGUAGE MODELS IN FIRE SAFETY PLANNING

Ori Ashkenazi & Shabtai Isaac

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

Alberto Giretti, Alessandro Carbonari &, Dilan Durmus

Polytechnic University of Marche, DICEA Department, via Brecce Bianche 12, 60131, Ancona, Italy

ABSTRACT: *This paper discusses the potential use of AI in general, and large language models (LLMs) in particular, to support knowledge management (KM) in the building industry. The application of conventional methods and tools for KM in the building industry is currently limited due to the large variability of buildings, and the industry's fragmentation. Instead, relatively labor-intensive methods need to be employed to curate the knowledge gained in previous projects and make it accessible for use in future projects. The recent development of LLMs has the potential to develop new approaches to KM in the building industry. These may include querying a variety of relatively unstructured documents from previous projects and other textual sources of technical expertise, processing these data to create knowledge, identifying patterns, and storing knowledge for future use. A proposed framework is defined for the use of LLMs for KM in construction. We will perform preliminary analyses on how to train models that can generate information and knowledge required to make decisions in the development of specific tasks of fire safety planning.*

KEYWORDS: *Large Language Models (LLMs), Knowledge Management (KM), Fire Safety Planning, Expert Systems (ESs), Artificial Intelligence (AI), Knowledge Graph, Ontology.*

1. INTRODUCTION

In the building industry, ensuring fire safety is of paramount importance. Effective fire safety planning plays a crucial role in mitigating risks, protecting occupants, and minimizing property damage (Kodur et al., 2019). However, the complex and dynamic nature of the building industry poses unique challenges in the realm of fire safety planning such as high fuel (fire) load, improper use of the materials, use of new construction materials with poor fire performance or longer response times for firefighting (Kodur et al., 2019; Parsamehr et al., 2023). Conventional methods and tools for knowledge management (KM) have struggled to adequately address these challenges due to the industry's inherent variability and fragmentation (Nikolic & Dakic, 2015).

Traditionally, fire safety planning has been based on expert knowledge (Maiellaro, 1997; Law & Spinardi, 2021). To curate and leverage the knowledge gained from previous projects for future endeavors, the building industry has relied on labor-intensive approaches. These approaches involve manually extracting and organizing information from disparate sources, making it time-consuming and resource-intensive (Liu, 1995). Consequently, the development and implementation of efficient fire safety planning strategies are hindered, hampering the industry's overall progress.

Expert systems (ESs) have been developed to capture and codify this expert knowledge, making it more accessible to fire safety professionals. However, ESs have a number of limitations, including the fact that they are difficult to maintain and update, and they can be inflexible in dealing with new or unforeseen situations (Tofiło et al., 2013).

Nevertheless, recent advancements in the field of Artificial Intelligence (AI), particularly in large language models (LLMs), offer new opportunities for overcoming knowledge management issues in the building industry. LLMs are a type of artificial intelligence (AI) that are trained on massive datasets of text and code (Shanahan, 2023). This allows them to learn the relationships between concepts and to generate text that is both informative and comprehensive. LLMs, such as OpenAI's GPT-3.5, have the potential to revolutionize the way knowledge is accessed, processed, and utilized.

The utilization of LLMs in fire safety planning offers the potential benefits of accessing and processing extensive data from previous projects, identifying patterns and trends in fire safety data, generating novel knowledge and insights, and adapting to new or unforeseen situations. These advancements can significantly improve decision-making processes and enhance the overall effectiveness of fire safety planning in the building industry.

Therefore, the primary goal of this paper is to investigate and assess the potential of LLMs in the context of fire safety planning. In pursuit of this objective, the current state of the art in expert systems (ESs) employed for fire

Referee List (DOI: 10.36253/fup_referee_list)
FUP Best Practice in Scholarly Publishing (DOI 10.36253/fup_best_practice)

Ori Ashkenazi, Shabtai Isaac, Alberto Giretti, Alessandro Carbonari, Dilan Durmus, *Transforming Building Industry Knowledge Management: A Study on the Role of Large Language Models in Fire Safety Planning,* pp. 729-738, DOI 10.36253/979-12-215-0289-3.73

safety planning will be conducted and LLMs will be introduced as an innovative approach to help define scenarios to ESs. Finally, the research will demonstrate the preliminary test results and conclude with future research directions in order to realize the full potential of LLMs for fire safety planning.

2. THE CHARACTERISTICS OF ESS AND LARGE LANGUAGE MODELS (LLMS)

AI systems process and analyze large datasets with the view of identifying patterns, relationships, drawing inferences, recommendations and taking action. With the advancement in AI, conversational AI came of age in 2010 to deal with the application of Natural Language Processing (NLP) to enable computers to interact with humans in a conversational way using natural language. The majority of the developed conversational AI agents in the AEC industry are based on the traditional approach to NLP, which requires time for processing the data, and users' interactions are often restricted as the agents are developed with the assumption of happy path users (Saka A. et al., 2023). Large Language Models (LLMs) are neural networks with large parameters and are trained using self-supervised learning and semi-supervised learning on large datasets. LLMs have improved NLP and shifted the direction away from training with labeled data for defined objectives. Generative Pre-trained Transformer (GPT) models which are decoder blocks only from OpenAI have gained significant attention and showed improved performance from GPT-2 (trained with 10 billion tokens) until the latest GPT-4 released in 2023. GPT models use transformer-based models that learn statistical patterns of natural language, enabling them to generate human-like language. One of the main advantages of GPT models is their capacity to produce language that is cohesive, fluent, and nearly indistinguishable from any text produced by humans. These models have been effectively used in a variety of applications, including chatbots, content generation, and machine translation. They can produce answers to open-ended questions, making them an important tool for natural language communication.

Not only have the communication and inference abilities demonstrated to emerge naturally in large language models, but even dedicated experiments have shown that chain-of-thought (COT) prompting improves performance on a range of arithmetic, commonsense, and symbolic reasoning tasks. Indeed, one's own thought process when solving a complicated reasoning task such as a multi-step math word problem entails decomposing that problem into intermediate steps and solving each of them before giving the final answer. A research endowed language models with the ability to generate a chain of thought, i.e. a coherent series of intermediate reasoning steps that lead to the final answer for a problem. It was reported that sufficiently large language models can generate chains of thought if demonstrations of chain-of-thought reasoning are provided in the examples for few-shot prompting (Wei J. et al., 2022). The motivations that led to the development of the chain-of-thought prompting method match the main goals that must be pursued in the field of knowledge management in construction. First, the chain of thought, in principle, allows models to decompose multi-step problems into intermediate steps, which means that additional computation can be allocated to problems that require more reasoning steps. Second, a chain of thought provides an interpretable window into the behavior of the model, suggesting how it might have arrived at a particular answer and providing opportunities to debug where the reasoning path went wrong. Third, chain-of-thought reasoning can be used for tasks such as math word problems, commonsense reasoning, and symbolic manipulation, and is potentially applicable (at least in principle) to any task that humans can solve via language. Finally, chain-of-thought reasoning can be readily elicited in sufficiently large off-the-shelf language models simply by including examples of chain-of-thought sequences into the examples of few-shot prompting (Wei J. et al., 2022). Among the main findings from this study, we would like to stress that chain-of-thought prompting does not positively impact performance for small models, rather it has larger performance gains for more-complicated problems. In addition, chain-of-thought prompting via GPT-3 compares favorably to prior state of the art, which typically finetunes a task-specific model on a labeled training dataset.

Although the construction industry is information-intensive and relies on myriad and diverse information from different stakeholders, there is a lack of information integration, reuse, and efficient management, all of which have an impact on the productivity of the industry. In other words, we must find out how to represent large amounts of diverse knowledge in a fashion that permits their effective use and interaction (Goldstein I, & Papert S., 1977). An expert system is a computer program whose performance is guided by specific, expert knowledge in solving problems. The problem-solving focus is crucial in this characterization because there must be the knowledge of central interest that can guide the search for solutions. The word "expert" implies narrow specialization (or focus) and substantial competence. It is intended to solve problems that are otherwise solved by people having dedicated training and exceptional skills. Thus, the standard of performance for expert systems is in human terms, by comparison with people carrying out a particular kind of task (Stefik M., 2014). Nowadays, expert systems are being used in a wide range of different interactive roles, such as smart spreadsheets, financial advisors, planning assistants, and cognitive coprocessors. In whatever role we employ expert systems, those systems require knowledge to be competent, the so-called knowledge base. Such knowledge must be elicited through effective techniques from multiple human experts. Then, techniques to process such knowledge should be adopted to work out a robust, reliable, and flexible expert system (Fekri-Ershad, S., 2013). Even in the most successful applications where expert systems outperform human experts in their reliability and consistency of results, expert systems have less breadth and flexibility than human experts. In this context, LLMs can help enhance the flexibility and reliability of expert systems, thanks to their ability to manage large amounts of information, even if it is not structured such as written documents, technical standards and regulations. GPT models are trained using vast amounts of unstructured text data, enabling them to generate language almost indistinguishable from human-generated text (OpenAI, 2023). The synoptic chart included below as Table 1, discusses how LLMs can be a potential tool for improving knowledge management in the building industry. LLMs can help widen the scope of expert systems thanks to their ability to scrape information from any (even) unstructured knowledge sources as soon as it is made available. Hence, they can increase and update the knowledge base of an expert system and set up decisional processes that integrate advanced paradigms.

Table 1: Leveraging LLMs to enhance decision-making in combination with expert systems

	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*

3. DISCUSSION OF GAPS AND COMPLEMENTARY ASPECTS

Few research works have applied GPT models in construction case studies. They mainly are concerned with information retrieval from BIM models and scheduling and sequencing tasks. These investigations helped find out inherent limitations, and they paved the way to an extended research review, that investigated opportunities and challenges about the practicality of GPT models in the construction industry (Saka A., 2023). These opportunities are categorized into different phases of the construction project lifecycle: predesign, design, construction, operation, and demolition.

In the pre-design phase, GPT models incorporated into predesign processes can help simplify decision-making, enhance communication among stakeholders, and speed up the discovery of design restrictions and possibilities. Similarly, they may support decisions with their powerful natural language processing and machine learning capabilities of evaluating data, delivering insightful suggestions, and promoting more efficient cooperation among stakeholders through real-time information and predictive analysis.

In the design phase, GPT models can automate regulatory compliance checks, reducing errors, and streamlining the design process. Other representative applications could be quantity take-off and costing (because GPT models can be leveraged by providing textual data for the model with necessary cost databases and estimation methods). If provided with information on standards, regulations, passive design principles, and renewable energy systems, GPT models can be leveraged for improving energy efficiency analysis.

In the construction phase, GPT models can capture and interpret textual information, enabling a more comprehensive understanding of construction project scheduling and logistics by means of a more flexible and intuitive approach than mathematical modeling. Other applications include dynamic risk identification and assessment by scrutinizing large volumes of project documents and historical data; enhancing progress monitoring and reporting by analyzing textual project updates and reports; site safety management; resource allocation and optimization; planning and organizing inspection and testing activities; claim and dispute resolution.

In the operation and maintenance phase, GPT can assist in predicting energy demand patterns, allowing for better planning and allocation of energy resources and providing customized recommendations. They could predict an asset's remaining useful life, predicting maintenance and repairs. They can be trained with relevant compliance documents, guidelines, and reporting requirements to capture patterns and embedded knowledge. These models can be used in compliance evaluation of facility management activities. Other applications might be space allocation for the best usage in a facility; generating sustainability reports; analyzing waste-related data.

In the demolition phase, GPT models can process large volumes of data, in order to: determine optimal demolition sequences, and recommend appropriate safety measures; optimize waste sorting and identify recyclable materials; offer the prospect of redeveloping and repurposing the site; analyze various risk factors and provide more accurate and objective assessments by leveraging their advanced natural language processing (NLP) and machine learning (ML) capabilities.

The analysis performed above suggests that some shared challenges must be faced:

1. Models are prone to hallucination, i.e. they give sound and plausible information that is not true, which reduces system performance and users' expectations.
2. Mainly structured data are usually needed in the fine-tuning of GPT models which are often not readily available in the construction industry; besides, availability and quality of data have been a major challenge for the application of artificial intelligence in the industry.
3. Although GPT models are large language models and trained on large data sets, their ability to understand domain-specific knowledge is limited. As such, there is a need for adequate fine-tuning of GPT models and the provision of context to improve their performance in a technical domain like the construction industry. Similarly, the regulations in the construction industry are many and vary over time.
4. GPT models are trained on large datasets. In the construction industry, data such as project design, cost, contracts, and schedules could be used as inputs, raising the concern of confidentiality of GPT models generating output with sensitive information.
5. There is often resistance to change in the industry. With the growing application of AI, industry practitioners and stakeholders are skeptical about trusting and accepting it.
6. Liability and accountability challenges are due to the training on a large amount of data. Bias, incomplete information and inaccuracies in the training data would affect the output generated by the models.
7. Deployment of GPT models in the construction industry requires new skill sets called "prompt engineering" and training programs for professionals. There are several techniques available for prompt engineering such as zero-shot (i.e. no examples are provided for the models to perform specified tasks), few-shot (i.e. provide contexts and examples in the prompts for the model) and chain-of-thought (CoT), i.e. prompting includes a series of intermediate reasoning steps.

In this paper, preliminary experiments about fine-tuning of GPT models for construction applications will be showcased. One of the objectives is to assess whether fine-tuned GPT models can accomplish a range of tasks with greater accuracy. Pre-training may involve unsupervised learning without labels or annotations. After pre-training, the model can be adjusted for a variety of tasks to increase the quality and accuracy of the text that is produced for that activity, such as language modeling, text categorization, or question-answering. Outcomes of our work include preliminary advice about how knowledge sources must be arranged and how queries can be prompted.

4. POTENTIAL FOR INTEGRATION

This paper discusses the hypothesis that the domain models on which ESs are based can be enriched and expanded through the analysis of relevant scenarios with the help of LLMs, by revealing variables and parameters that are relevant but are currently either excluded from the ES or incorrectly defined in it due to the modeler's bias. Four components are essential to apply such an approach: scenarios, a LLM, an ES and a validation process.

<u>Scenarios:</u> the first step is to identify relevant scenarios that represent real-world situations within the domain of the ES. In this case, the relevant domain is fire safety, and scenarios of fire events are therefore collected. These scenarios should cover a wide range of situations to ensure the system can handle various cases (i.e., different types of fire events). Of particular interest are scenarios that may not occur frequently, but that are still important to handle, as they may reveal aspects that will need to be incorporated into the model of the domain.

<u>LLMS:</u> LLMs can be used to identify patterns, trends, and correlations within the scenarios, thus enriching and

expanding the domain model. The process of applying a LLM should be iterative, so that the model can continuously be refined as new scenarios are identified or as the domain evolves. By querying the LLM, different aspects of a scenario can be explored. Each query can build on knowledge gained from the previous queries, allowing the analyst to gradually gain a deeper understanding of the domain. LLMs can challenge existing assumptions and preconceptions, by revealing factors that might have been previously overlooked. They may thus help in avoiding confirmation bias and encourage a more open-minded analysis. Scenarios involving complex interactions between different parameters might be better understood by querying the LLM and exploring how the different factors influence each other. In addition, the analyst can vary certain parameters in the queries to understand how the changes might impact the scenario's outcomes, performing what may be comparable to a sensitivity analysis.

<u>ESs:</u> By repeatedly engaging with a LLM through queries to analyze specific scenarios, the analyst may uncover certain aspects that are otherwise ignored, or that are not well understood. By incorporating those aspects in an existing structured knowledge graph of the ES, the accuracy of the recommendations provided by the ES for such scenarios is increased. This is done by converting the information that is gained into explicit logical statements or rule sets to address varying situations, in a way that represents the real-life decision-making process. The logical statements should be structured in a format suitable for their integration into the existing knowledge graph of the expert system, for example by creating new data fields or modifying existing ones to accommodate the additional information. To account for complex relationships between data in the knowledge graphs, hierarchies of rules can be established, where higher-level rules encompass general scenarios, and lower-level rules handle exceptions and specific cases. This hierarchical structure may support a more detailed and accurate decision-making process.

<u>Validation:</u> The updated knowledge graph of the ES, and the logical statements on which it is based, should be validated by consulting with domain experts, to ensure their accuracy and completeness. Once the enriched knowledge graph is integrated into the existing ES, the system should be thoroughly tested with both previously used and new scenarios to ensure that it performs accurately and reliably. The iterative nature of the process of continuous improvement ensures that the system stays up-to-date and relevant as it learns from new data and insights over time.

To summarize, by enriching and expanding the knowledge graph of an ES through the analysis of relevant scenarios with a LLM, the system can become more robust and capable of supporting informed decision-making in complex real-world situations. By incorporating the aspects uncovered through the querying of a LLM, and defining relevant logical statements, the ES can improve the accuracy of its recommendations for a broader range of scenarios.

5. PROPOSED APPROACH FOR THE USE OF LLMS TO HELP DEFINE SCENARIOS FOR ESS

In the first stage, a LLM is queried that has embedded documents describing scenarios. In GPT (Generative Pre-trained Transformer) models, relevant documents can be embedded to help the model better understand context, and thus improve the accuracy of answers to queries. The embedded documents can be used to fine-tune the GPT model for a specific task. By linking the documents with a query in a single prompt, a unified input is created for the model, allowing it to use the context from the documents to provide more accurate answers to the given queries. Naturally, the quality and relevance of the documents used in the prompt will play a significant role in determining the accuracy of the answers provided, and their selection is therefore crucial.

In the second stage, the answers provided by the LLM are used to enhance the ES. This can be done directly in the ES's knowledge graph, or indirectly through an ontology on which the knowledge graph is based (Figure 1).

The ES's knowledge graphs can be enhanced directly with a set of knowledge units gained from the scenarios by querying the LLM. To do so, the LLMs answers, which are provided in the form of natural language text, must be converted into a structured format that is suitable for their incorporation into the knowledge graph. Typically, entities, relationships, and attributes are extracted from the text, and converted into corresponding nodes, edges, and properties in the knowledge graph, while maintaining the graph's integrity. Depending on the knowledge units that are extracted, existing nodes in the knowledge graph may need to be updated or merged, or new nodes added. Preferably, the newly added knowledge units are linked with the relevant scenarios on which they are based, to maintain traceability and support additional future updates.

Alternatively, the ES's knowledge graph can be enhanced indirectly through an ontology on which it is based. To do so, a semantic ontology needs to be defined and used to create a conceptual bridge between the two systems. This ontology is on the one hand used to define the ES's knowledge graph and is on the other hand repeatedly updated based on the information gained through the interaction with the LLM. The knowledge units, acquired by querying the LLM, are used to define new classes, properties, and relationships in the ontology, thus extending the ontology to better represent the domain. Following this, the ES's knowledge graph is updated based on the extended ontology. Preferably, version control is implemented for the ontology to support the tracking of changes and management of updates. By repeatedly applying the previously described process, the ES can become over time capable of handling a broader range of scenarios with greater accuracy and relevance.

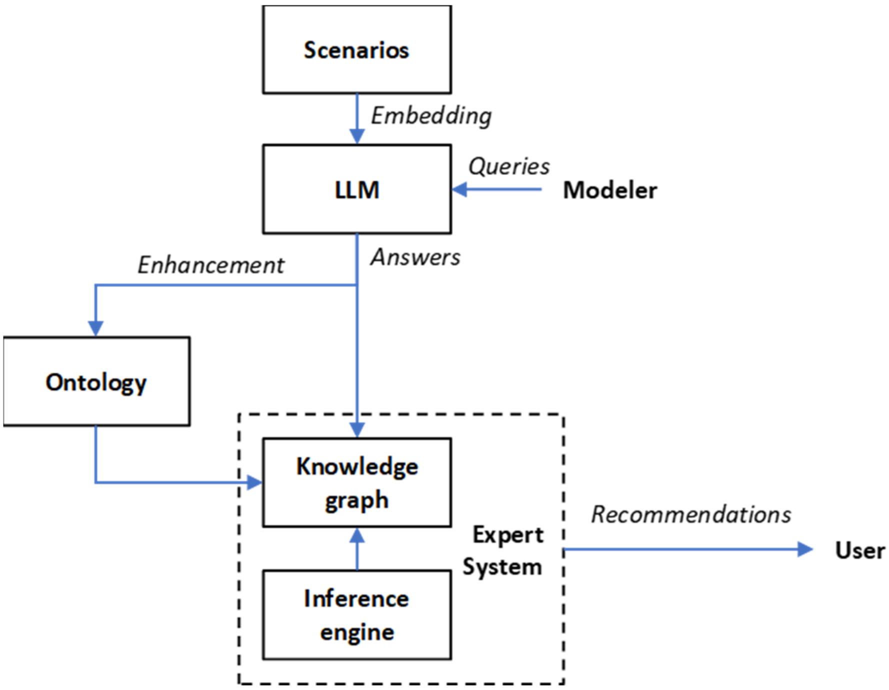

Figure 1: Enhancement of an Expert System using a Large Language Model

6. CHOICE OF THE APPLICATION DOMAIN

The application domain of fire safety planning was chosen for its paramount significance in the building industry and the intricate challenges it poses in safeguarding human life and property from the devastating impacts of fire incidents. Fire incidents in buildings can lead to devastating consequences, including injuries, fatalities, and significant property damage.

By leveraging Large Language Models (LLMs) to systematically analyze reports of real-life fire events, pertinent insights into the factors influencing fire incidents and the efficacy of existing fire safety measures can be extracted. The LLMs provide natural language responses, which can be converted into structured knowledge units suitable for integration into the Expert Systems (ESs) knowledge graph.

Querying reports of fire events and subsequently expanding the ontology based on these queries aligns coherently with the standards set forth by the National Fire Protection Association (NFPA), particularly in Chapter 5, clause 5.5.3. (National Fire Protection Association, 2017). The NFPA emphasizes the criticality of selecting pertinent scenarios, both from predefined options and those deemed relevant by designers, to ensure a comprehensive approach to fire safety planning.

The process of querying, extracting insights, and expanding the ontology facilitates the inclusion of diverse scenarios, encompassing not only common cases but also rare and significant events. The adherence to NFPA standards ensures that the fire safety planning process remains thorough and adaptive, contributing to informed decision-making and improved fire safety strategies in building environments. Consequently, this integration demonstrates the potential to advance knowledge management in the fire safety domain, elevating the overall efficacy of fire safety planning measures.

In the upcoming chapter, a series of tests will be conducted using selected case studies as examples such as, NFPA case study of nightclub fires (Duval, R.F., 2006):

- Rhythm Club, Natchez/Mississippi
- Cocoanut Grove, Boston/ Massachusetts

These case studies, derived from published reports, have been carefully chosen to represent diverse fire safety scenarios encountered in real-world building environments. The adherence to NFPA standards ensures the appropriateness and validity of utilizing published reports for the case studies. These tests will provide valuable insights into the efficacy of the proposed approach, showcasing its potential to enhance knowledge management in fire safety planning and inform more informed decision-making processes.

7. IMPLEMENTATION

To expand the relevant ontology, a modular and contextually aware approach is devised to querying a knowledge corpus for "low temperature" knowledge of a predetermined format, as represented in the ontology. For this, the following components are used.

7.1 Agent-based Methodology

To support the querying of the said corpus, agents will be created and utilized. When creating an LLM-based agent, understanding the nature of the action, or "step" executed by the agent is paramount. The steps can be categorized into:

- User-guided: A user aligns the agent with the target by adding specific information.
- Pre-prompted: The system operates with a generic instruction set to achieve the desired outcome.
- Autonomous: The agent compiles a list of tasks before engaging in user-guided or pre-prompted tasks.

Each of these steps necessitates a mechanism for handling context, as all prompts require appropriate context to yield optimal results. The context influences the stochastic and statistical process of token generation. In this article, we have focused on one step, node querying, which is a pre-prompted step. All steps must be accompanied by the right context. The essence of context handling lies in the creation, storage, and injection of the right bit of context in the proper format (length, form, keywords, wrapping), providing an agent with a solid foundation for action. This mechanism embodies an object-oriented, modular scheme where pre-prompts and context are treated as retrievable objects for injection. These objects are stored in a tree graph, and their predefined format can be utilized for stable injection. Context objects must be designed according to the types of steps and the theme of the project. Each theme presents different objects and relationships, demanding specific syntax to harness hidden semantic information.

The creation of context, being paramount, can be realized through two processes:

- Demanding a Specific Output Format: As part of the agent's response, a certain output format is required for the context object.
- External Ingestion of Information: Utilizing an external mechanism to ingest existing information and represent it in a context object.

As the methodology suggests, the step presented here could be used to continuously enlarge the scope of querying or enable working in parallel on other tasks while offering the right context for them.

7.2 Prompting

To establish the correct behavior for the agent, a system prompt is created first, based on the ontology description and a leading "persona" statement (a role with which the agent's behavior will be aligned).

Example:

```
You are a Fire Evacuation Specialist working with an ontology.

The evacuation of buildings in fire emergency situations is a problem for
which solutions must be found that can help the occupants of these spaces,
guiding them along their route until they are safe. The purpose of the
ontology is to build a knowledge Graph that can contribute to a better
understanding of this topic, as well as to the development of solutions or
systems for the evacuation of buildings more capable of guiding the
occupants of these spaces to a safe place.

Ontology Domain and Scope are:
```

```
Domain: Evacuation of buildings under fire emergency

Focus: Recommendation of evacuation routes in real time based on contextual
information obtained from loT devices.

Contribution: Contributes to a solution for the evacuation of buildings
under fire.

Usage: In fire building evacuation systems.

Purpose: Support interoperability between IoT devices, occupants and other
systems.
```

Next, a ontology python class is created, which is an object-oriented, programming language-resembling representation of an ontology class. The instance can have a description generated for it using the `str` function.

Example:

```
data_property1 = DataProperty("Class_Name", "Property_Name1",
"Description1")

ontology_node = OntologyClass(name="Class_Name", comment="Comment about the
class.", dataproperties=[data_property1, data_property2])
```

The attributes are described as:

```
f"The attribute {self.name} (an attribute to {self.name}) asks:
{self.question}"
```

And the class is then described as:

```
f"The class {self.name} is (are) {self.comment}\nFor finding information to
fill in attributes, consider the following:\n{dataproperties_str}"
```

This representation is used as the "query_node_description".

Next, an example of the "output_format" we demand the agent to follow is given:

```
Formulate your response based on the rule:

1. If you don't know the answer, just respond "Don't know". DO NOT try to
make up an answer.

2. If the question is not related to the context, respond that you are
tuned to only answer questions that are related to the context.

3. DO NOT add anything else to the response.

4. Be clear and precise.

5. Fill the template with information from the context, fill all the
requested data types with information.

Response generic format:

Name: <name>, Description: <description>, Data properties: [<data
properties>]
```

This information is then accompanied by context queried from the knowledge corpus. Texts are extracted from a PDF file and converted into numerical vectors using a specialized PDF processor. This vector representation of the text is stored and then queried using the node description for the retrieval of the top-k results - the context holding the "high-temperature" information.

The augmented prompt is constructed using this formula:

```
Use the following pieces of context to formulate an ontology class instance
representation.

Context:

{context}, {query_node_description}, {output_format}
```

The augmented prompt is then sent to the GPT-4 API for processing and lowering the information's temperature into the required structure.

The above described process will result in the successful curation of information, in a desired structure of one class instance, as it is portrayed in the text. The following are key aspects of the process:

- Data Collection: Utilizing the GPT API to get data from a knowledge corpus to aggregate and accommodate data.
- Knowledge Base Enrichment: Briefly describing the methodology to enrich the knowledge base, including the integration of ontologies, vector handling, and prompt engineering.
- Class Instance Formation: Achieving the structured representation of one class instance from the text, adhering to the defined ontology and prompt templates.

8. CONCLUSIONS

In conclusion, this paper discusses the transformative potential of Large Language Models (LLMs) in Knowledge Management (KM) within the building industry, with a particular focus on fire safety planning. By harnessing the computational prowess of LLMs, a framework is delineated capable of mining unstructured documents from previous projects, processing this data into actionable knowledge, and preserving it for future endeavors.

The challenges for future development of such a framework are multifaceted, and include:

- Accuracy Benchmarking: Creating a mechanism to assess the correctness of the curated information is crucial for ensuring the agent's trustworthiness. This involves evaluating the relevance and consistency of the curated low-temperature information, and employing a back-propagation process to analyze its alignment with the specific domain, ontology, and templates.
- Step Size in Respect of Token Attention Span: The token attention span of the model must be meticulously managed to ensure that the curated information is accurate and properly compiled when augmented into the prompt template. This encompasses careful handling of tokenization, embedding, prompt consciousness, and robust mechanisms for large ontology structures and knowledge corpus.
- Fractured Information Retrieval: Addressing the challenge of incomplete knowledge sources is vital to collect comprehensive information without resorting to premature termination by the model.
- Unwanted Event Amalgamation: Further refinement in the ingestion of PDF files is needed to prevent the creation of class instances that may not be accurately represented in the text.

The outlined challenges, such as accuracy benchmarking and token attention span, underscore the complexities inherent in this approach. Nevertheless, surmounting these hurdles may pave the way for transitioning from labor-intensive curation methods to an automated, intelligent system. This paradigm shift has the potential to change how knowledge is managed in an industry characterized by fragmentation and variability, offering a more streamlined and effective method for capitalizing on past insights to fuel future innovation.

REFERENCES

Duval, R.F., 2006. NFPA Case Study: nightclub fires. National Fire Protection Association, Fire Investigations Department.

Fekri-Ershad, S., Tajalizadeh, H., & Jafari, S. (2013). Design and Development of an Expert System to Help Head of University Departments. arXiv preprint arXiv:1308.0356.

Goldstein, I., & Papert, S. (1977). Artificial intelligence, language, and the study of knowledge. Cognitive science, 1(1), 84-123.

Kodur, V., Kumar, P., & Rafi, M. M. (2019). Fire hazard in buildings: Review, assessment and strategies for improving fire safety. PSU Research Review, 4(1), 1–23. https://doi.org/10.1108/PRR-12-2018-0033

Law, A., & Spinardi, G. (2021). Performing Expertise in Building Regulation: ‘Codespeak’ and Fire Safety Experts. Minerva, 59(4), 515–538. https://doi.org/10.1007/s11024-021-09446-5

Liu, L. Y. (1995). Digital data-collection device for construction site documentation: *Proceedings of the 2nd Congress on Computing in Civil Engineering. Part 1 (of 2). Computing in Civil Engineering (New York)*, 2, 1287–1293.

Maiellaro, N. (1997). Using expert systems to check building applications. 19.

Nikolic, B., & Dakic, J. (2015). Knowledge management in the function of risk assessment. 3(2).

National Fire Protection Association (Ed.). (2017). NFPA 101: Life safety code 2018 (2018 edition). National Fire Protection Association.

OpenAI. (2023). Fine-tuning Retrieved 2nd of May 2023 from https://platform.openai.com/docs/guides/fine-tuning/fine-tuning

Parsamehr, M., Perera, U. S., Dodanwala, T. C., Perera, P., & Ruparathna, R. (2023). A review of construction management challenges and BIM-based solutions: Perspectives from the schedule, cost, quality, and safety management. Asian Journal of Civil Engineering, 24(1), 353–389. https://doi.org/10.1007/s42107-022-00501-4

Saka, A., Taiwo, R., Saka, N., Salami, B., Ajayi, S., Akande, K., & Kazemi, H. (2023). GPT Models in Construction Industry: Opportunities, Limitations, and a Use Case Validation. arXiv preprint arXiv:2305.18997.

Shanahan, M. (2023). Talking About Large Language Models (arXiv:2212.03551). arXiv. https://doi.org/10.48550/arXiv.2212.03551

Stefik, M. (2014). *Introduction to knowledge systems.* Elsevier.

Tofiło, P., Konecki, M., Gałaj, J., Jaskółowski, W., Tuśnio, N., & Cisek, M. (2013). Expert System for Building Fire Safety Analysis and Risk Assessment. Procedia Engineering, 57, 1156–1165. https://doi.org/10.1016/j.proeng.2013.04.146

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems, 35, 24824-24837.*