



TOWARDS AUTOMATION IN COST ESTIMATION: LLM-BASED METHODOLOGY FOR CLASSIFYING AND EXTRACTING COST DATA FROM TENDER DOCUMENTS

Chiara Gatto¹, Claudio Mirarchi¹, and Alberto Pavan¹

¹Politecnico di Milano, Milano, IT

Abstract

Cost estimation in building industry largely relies on manually extracting and classifying textual descriptions, a process susceptible to human error. Although recent advancements in Large Language Models (LLMs) hold promise, their application in this domain requires further investigation.

This study proposes a methodology to optimize LLM performance validated through the development of a tool that classifies cost descriptions into a three-level hierarchical taxonomy and extracts relevant information organising the data in a database as output. Results demonstrate a F1 score of 0.96 on classification tasks contributing to cost estimation automation, reducing manual processing, and enhancing knowledge management within the domain.

Introduction

Accurately estimating costs is a fundamental aspect of construction project planning, particularly during the tendering phase. This process involves forecasting expenses, timelines, and resource allocations to ensure both financial feasibility and effective project execution. A well-structured cost estimation not only supports investment planning but also enhances decision-making throughout the construction phase, reducing financial risks and improving overall project efficiency (Sepasgozar et al., 2022).

The traditional approach to cost management relies heavily on manually inspecting textual documents to extract information related to various factors, such as the type of materials used, construction techniques, site conditions, project duration, and the impact of regulations and risks. In public procurement, selecting a cost item for an estimate is a critical step in the contractual process, as it becomes an integral part of the contract. If discrepancies or issues arise later regarding this item, the entire process could lead to legal disputes. This makes it essential to ensure precision and clarity in selecting and defining cost items to avoid conflicts during the project execution (Giuda et al., 2020). Therefore this process requires a high level of expertise and is subject to two major challenges. First, it is time-consuming and labor-intensive, leading to inefficiency. Second, the process is influenced by subjective factors, such as the emotions and fatigue of the personnel in-

involved. These variables can introduce inconsistencies and errors, further complicating the already demanding task of analysing construction cost related documents (Lee et al., 2019).

Generative AI, particularly Large Language Models (LLMs), has revolutionized human-computer interaction by enabling advanced natural language processing. Trained on vast datasets, these models excel in language comprehension and generation, quickly adapting to tasks with minimal instruction. This allows them to handle unseen data and outperform traditional models. Reduced technical investment costs have further expanded their applications, in risk management, compliance checking, contract review with including task such as information extraction, summarization, reasoning (Wan et al., 2024).

However, their application in the cost sector is still in its early stages, and the assessment of their full potential has only just begun. In AECO cost estimation, LLMs face difficulties in handling complex, nested, and conditional regulatory statements, posing challenges for broader industry adoption. Additionally, limited research has explored the use of APIs to scale testing across larger datasets, leaving a gap in understanding their scalability and effectiveness in real-world scenarios (Chen et al., 2024).

This study explores the potential of LLMs in cost management within the construction sector. Specifically, it proposes a methodology to enhance the accuracy of classify work item descriptions according to a three-level knowledge hierarchical system. This approach allows the model to progressively adapt to specific contexts, facilitating a subsequent phase of relevant information extraction. The study introduces a tool architecture based on dynamic prompts and an AI-driven pipeline that leverages OpenAI GPT 4.0 APIs to develop a tool capable of automating the classification and extraction of data from cost descriptions. This paper is structured as follows: the "Background" section provides an overview of the state of the art in the field. The methodology developed is presented in the following section, which is then applied in the "Application to a Practical Case Study," where both the case study and its implementation are detailed. The effectiveness of the methodology is evaluated in the "Results" section. Finally, the paper concludes with a discussion of the findings and potential future developments.

Background

Effective knowledge representation is essential for enhancing data accessibility, decision-making, and communication among project stakeholders. It ensures efficient data retrieval, facilitates structured decision-making, and minimizes errors by aligning all participants with a shared understanding of project specifications (Taye, 2010). Furthermore, organizing information systematically supports the integration of digital technologies and ensures adherence to national regulations. A well-structured approach to data management improves accuracy and quality in construction projects, making it a critical component of successful project execution (Pupeikis et al., 2022).

To achieve this structured representation, ontologies play a crucial role by formally defining the concepts, entities, and relationships within the construction domain (Gruber, 1993).

Beyond ontological representation, various classification systems have been developed to structure AECO-related information and meet diverse user needs. For instance, Uniclass 2015 provides a unified system applicable across all construction sectors in the UK and beyond, while OmniClass serves as another widely recognized standard. These approaches, often complemented by ontological models, contribute to efficient project delivery and foster international collaboration.

In addition to industry-wide classification standards, internal project management structures rely on classification methods such as the Work Breakdown Structure (WBS). This is a hierarchical system that decomposes complex construction projects into smaller, manageable tasks, ensuring clarity and organization. This structured breakdown facilitates accurate cost estimation, resource allocation, and scheduling while reducing errors and improving decision-making. By supporting task delegation and monitoring, WBS enhances productivity and safety, ensuring smooth project execution throughout its lifecycle (Stas and Abrishami, 2024).

In the cost domain, limited research has focused on the automatic structuring of information. One notable study by Martínez-Rojas et al. (2013) presents a preliminary approach for the automatic classification of Work Descriptions in construction projects, regardless of their linguistic framework or original structure. The proposed methodology enables the categorization of work descriptions from various sources, improving consistency and accessibility of project-related information.

Building on this initial work, Martínez-Rojas et al. (2018) expanded the study by testing the approach on a larger dataset of work descriptions. Their research explored six different classification techniques—C4.5 decision tree, random forest, Naïve Bayes, neural networks, support vector machines, and k-nearest neighbours to assign work descriptions to a standardized hierarchical structure of task groups. To enhance classification accuracy, the dataset underwent basic linguistic preprocessing, including text cleaning and synonym replacement, reducing vocabulary

complexity. The study demonstrated high classification accuracy across a broad range of work description categories, highlighting the effectiveness of automated classification techniques in structuring construction-related data. A subsequent study by Moon et al. (2021) developed a Bi-LSTM-based Named Entity Recognition (NER) model to extract information from construction specifications, classifying data into five categories: Organization, Action, Element, Standard, and Reference. This study marks one of the first successful applications of NER in the construction industry, showcasing its potential for automating information extraction.

A recent study presents a methodology for automating the processing of work descriptions, laying the groundwork for automated Quantity Take-off (QTO) and cost estimation through an NLP-based information extraction model that integrates a Hidden Markov Model (HMM) and formalized labeling rules. However, the study's scope is limited as it relies on RSMeans cost items, which may not be applicable to cost items from other sources. Additionally, the dataset used for training and testing is limited Tang et al. (2022).

Wang et al. (2021) proposes a multi-scale information retrieval approach for BIM, combining its hierarchical structure with Natural Language Processing (NLP). The aim is to develop a method for locating building components or attributes based on unified queries. The BIM hierarchy is represented through a BIM Hierarchy Tree (BIH-Tree) model, while NLP and the International Framework for Dictionaries (IFD) are used to parse and standardize the queries.

The studies presented utilize NLP techniques, a subset of AI that helps machines understand human language by analysing text structures and word meanings. NLP is increasingly adopted in the AEC sector for tasks such as information extraction, document organization, expert systems, and automated compliance checking (Wu et al., 2022).

Within this context, generative AI, particularly LLMs, has revolutionized natural language processing. These models adapt quickly to new tasks with minimal input, outperforming traditional models, and are increasingly used in risk management, compliance checking, and contract review, including tasks like extraction, summarization, and reasoning (Wan et al., 2024).

In the context of construction projects, LLMs have been tested for generating schedules for simple tasks. The authors proposed a method that uses as input project details such as work descriptions, materials, and deadlines, and set as output a task list and schedule. Practitioners evaluated the results, finding them coherent and in line with the project scope, even if tested over a limited dataset (?).

A further study focused on retrieving information from BIM models using natural language queries tested the effectiveness of the GPT language model. The authors developed a prompt generator to interpret users' natural language queries and a query management system to struc-

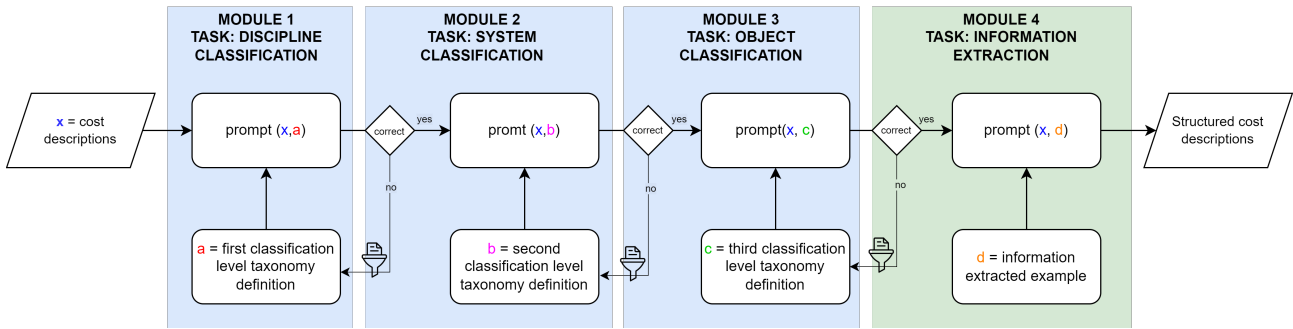


Figure 1: Methodology chart.

ture the queries for retrieving information from external databases (Zheng and Fischer, 2023).

In the context of compliance checking Chen et al. (2024) presents an automated framework for construction projects compliance checking that integrates LLMs, deep learning models, and ontology knowledge models. The proposed framework leverages the few-shot learning capabilities of LLMs to reduce the need for large annotated datasets and integrates deep learning for preliminary classification of regulatory texts. This innovative approach automates the processing of regulatory texts, enhancing both the efficiency and accuracy of compliance checks, making it a scalable solution for the construction industry .

Moving to the domain of building regulation, Fuchs et al. (2024) investigates using GPT-3.5 to automate converting building regulations into a computable format for automated compliance checking (ACC) in construction. By applying few-shot learning, GPT-3.5 learns the structure of LegalRuleML, translating regulations into machine-readable formats. The paper explores leveraging domain knowledge and techniques like chain-of-thought reasoning to improve accuracy, suggesting that LLMs can enhance the efficiency of compliance checking in construction.

Regarding the context of risk management, Ashkenazi et al. (2023) paper explores the potential of AI, specifically large language models (LLMs), to enhance knowledge management (KM) in the building industry. Due to the industry's fragmentation and the variability of buildings, conventional KM methods are labour-intensive and inefficient. LLMs offer the possibility of automating the process of querying unstructured documents, processing data, identifying patterns, and storing knowledge for future use. The paper proposes a framework for applying LLMs to KM in construction and outlines preliminary analyses on training models for decision-making in fire safety planning.

The state of the art shows a growing trend in using AI, particularly the study just cited promoted the use of LLMs to enhance knowledge management. While their potential as decision-support tools is clear, studies are still few in AECO sector and in early stages. Particularly, LLMs have yet to be tested in the cost estimation sector, a critical phase requiring meticulous planning of materials, time, and resources where the use of those techniques can optimize the

process.

Methodology

This section outlines the methodological approach adopted in the study, which is also summarized in Figure 1. The proposed methodology forms the basis for designing the dynamic prompt employed in the tool development pipeline, aimed at automating the processes of classification and information extraction.

As illustrated in the diagram, the methodology is structured around the two main tasks the tool is designed to perform:

- classifying the cost description, highlighted in the blue boxes;
- extracting entities from entities, highlighted in the green box.

Additionally, the model is tailored to the specific case study, which involves a three-level hierarchical classification system. The type of information extracted varies based on the cost item category; for example, data related to insulation layers differ from those concerning pavement materials.

A key aspect of this methodology is the organization and definition of domain-specific knowledge, which refers to how practitioners choose to structure the knowledge. This structure may vary in complexity, with knowledge organized across two, three, or more levels, depending on the specific requirements of the practitioner or the organization. The number of entities to be extracted can also differ, as it is influenced by the particular needs of each practitioner or company.

Each class within the knowledge structure must be clearly defined to ensure accurate classification and extraction. Since the model uses a multi-label classification approach, where each description can belong to multiple classes (one for each module), it is crucial to establish a precise taxonomy at each hierarchical level. This step ensures consistency across classifications and ultimately improves the model's performance.

The next step involves designing dynamic prompts that can handle variable input based on the context. These

prompts are dynamic in nature because the response generated by each one influences the construction of the subsequent prompt. This allows the model to process information progressively, starting from broader, more general input and gradually focusing on more detailed and specific aspects.

The prompts are chained one to the others, in a way that the output of the first affects the composition and the output of the following one. This staged approach ensures that the model is not overwhelmed by irrelevant information, avoiding to introduce details that don't relate to the label description identified in the previous classification step.

As the methodology progresses, it follows an iterative process. If the model fails to accurately assign the appropriate classes, the cycle returns to refining the category definitions. This step is crucial, as it helps the model recognize when the descriptions it receives are too general or ambiguous to be classified correctly.

Finally, the framework was tested by its practical implementation by using OpenAI API to recall GPT 4.0 model in Google Colab environment.

Application on a practical case study

Case study

The proposed model was applied and tested on a practical case study using 70 textual work item descriptions from the Lombardy Region Price List.

In accordance with the Italian Public Contracts Code (Art. 23, D.lgs 50/2016), each region is required to publish an annual price list that establishes the cost baseline for public tenders. These lists provide practitioners with unit prices for construction work items and a catalogue of elemental resources, ensuring pricing transparency. Regular updates introduce price adjustments, new entries, and the removal of obsolete items. Since price list descriptions are presented in natural language without a standardized format, their interpretation is inherently ambiguous, sometimes leading to legal disputes. The cost item definitions analysed in this study are the same as those investigated in previous papers addressing the challenge of parametrizing cost information from natural language descriptions of

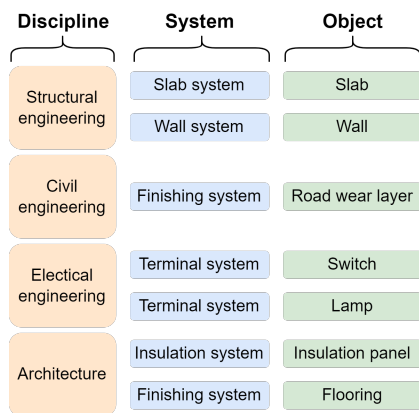


Figure 2: Case study.

price lists documents (Gatto et al., 2023; Cassandro et al., 2024).

To assess the model's adaptability to different types of work items, the dataset was evenly distributed across multiple categories, covering four disciplinary areas and seven object typologies as showed in figure 2: slab and wall belonging to the Discipline Structural Engineering; switch and lamp belonging to the discipline Electronic engineering; Road wear layer belonging to the civil engineering; Insulation panels and flooring belonging to the Architecture discipline.

Regarding the definition of domain knowledge, the methodology is designed to be adaptable to various contexts, following the classification logic of WBS. In this specific case, the hierarchical knowledge structure is based on a framework previously developed to efficiently organize large volumes of work item descriptions, tested using the Lombardy Region Price Lists (Cassandro et al., 2023). The hierarchical knowledge representation is divided into three levels: from the most general, relating to the discipline domain where the work item is applied, to the more specific systems domain that encompass various types of building elements composing a building system based on the function and type of application, and finally, the most specific object level, which identifies the type of object that is the subject of the work item.

Framework development

This section presents the implementation of the methodology, beginning with dataset preprocessing and the definition of knowledge representation.

First, the dataset was structured in JSON format to facilitate dynamic and organized management of information. During the dataset preparation, it was essential to assign definitions to each system class. The classification process followed a hierarchical structure, starting with the identification of disciplines at the first level, followed by system definitions, and finally, the characterization of individual objects.

As shown in Figure 3, the descriptions range from general definitions at the discipline level to increasingly detailed specifications of individual objects. However, the object definitions remain simple and not excessively detailed, ensuring flexibility to represent all categories that a building project may contain for a given object type.

Regarding the classification task, the model employs a

```

{
  "Discipline": "Architecture",
  "System": "Insulation system",
  "Object": "Insulation panel",
  "Discipline_definition": "Related to all those works of non-structural type performed on civil buildings",
  "System_definition": "Related to all those works which improve the thermal, acoustic and hygrothermal performance of a building, enhancing the protection of the indoor environment from external agents and improving its energy comfort"
},
{
  "Discipline": "Structural engineering",
  "System": "Slab system",
  "Object": "Slab",
  "Discipline_definition": "Related to all those works of structural type performed on civil buildings",
  "System_definition": "Related to all those works performed for the construction, maintenance or demolition of the structural parts constituting a slab"
}

```

Figure 3: Dataset preprocessing.

three-level prompt architecture, which executes one task per stage, as illustrated in Figure 4. The output of this first prompt influences the second one, where only the systems associated with the selected discipline are retained and passed along. The second prompt receives this filtered subset of systems and, in turn, its output determines the composition of the third prompt, which will only include the list of objects related to the selected system.

Thus, each level of prompt composition and content is shaped by the specific context of the previously identified elements. The process continues iteratively, progressively narrowing the classification scope until the final object is precisely identified.

This approach was adopted to increase the likelihood of obtaining accurate responses by directing the model to focus on few filtered classes rather than the entire set of possible categories. In this way, the risk of ambiguity and subtle differences between building objects (e.g., a thermal insulation panel in the Architectural discipline and a fire insulation panel in the Fire safety discipline) is reduced, pushing the model to operate in a more deterministic rather than probabilistic manner. Subsequently, the process undergoes an iterative refinement process where, if the classification result is unsatisfactory, the class definitions are revisited and clarified.

Once the description of the construction work has been classified, the next task implemented through prompt structuring is to determine the type of information to extract for each specific object. To guide this step, a few-shot prompting approach is adopted, as illustrated in Figure 4. For each object type, two examples of text de-structuring are prepared in JSON format. These examples demonstrate how to extract relevant information from the description and indicate the expected structure of the output. By providing these structured examples, the model is able to return processed, machine-readable data, enabling the automation of subsequent extraction steps, as shown in the module 4 stage of the Figure 4.

Results

This section presents the results of the methodology applied to the previously outlined case study. It is structured into two subchapters, each corresponding to a specific task performed by the LLM. All results have been manually verified to compute the relevant metrics. On average, the tool took 28 minutes to process the entire dataset. In comparison, a manual evaluation from a domain expert would take approximately 3 minutes per description, meaning a total of 210 minutes.

Classification task

Regarding the classification task, the metrics used to evaluate the results include the standard Precision, Recall, and F1 score, with the corresponding formulas provided below.

$$P = \frac{TP}{TP + FP} \quad (1)$$

$$R = \frac{TP}{TP + FN} \quad (2)$$

$$F1 = \frac{2 \cdot P \cdot R}{P + R} \quad (3)$$

Tables 3 and 2 present the results obtained from the first test and the second iteration, respectively. The initial application of the framework did not produce fully satisfactory results. The best performance was achieved in addressing discipline classes, while accuracy decreased progressively at the system and object levels. In particular, the model struggled to correctly assign the discipline for cost descriptions related to insulating panels. As a result, it consistently misclassified these elements at both the system and object levels. A similar issue emerged with descriptions of the road wear layer: while the model correctly identified the "Civil Engineering" discipline, it incorrectly assigned the system class in 9 out of 10 cases. This misclassification prevented the correct identification of the corresponding object in the final classification step.

Table 1: Classification task evaluation metrics on the 1st iteration.

Class	P	R	F1
Disciplines	0.83	0.87	0.82
Systems	1.0	0.66	0.70
Objects	0.0	0.65	0.69

After analysing the initial results, a refinement process was undertaken to improve classification accuracy by providing more precise definitions for the misclassified information classes. For instance, the definition of the Architecture discipline was revised from "Related to all those works of non-structural type performed on civil buildings" to "Related to all those works of non-structural type performed to enhance liveability and internal comfort in civil buildings."

This adjustment addressed a specific misclassification issue: previously, the model had incorrectly assigned Structural Engineering to insulating panels, likely due to its training data, which included structurally relevant insulating panels. With the refined definition, the model correctly associates these elements with the Architecture discipline, improving classification performance.

The same refinement process was applied to other misclassified categories, leading to a noticeable improvement in the framework's performance. By iteratively adjusting class definitions, the model was better able to distinguish between similar categories, resulting in more accurate classifications.

Information extraction task

The same metrics are used to evaluate the second task, considering True Positives (TP) as the number of correctly extracted fields, False Positives (FP) as the erroneously extracted fields, and False Negatives (FN) as the fields

Table 2: Classification task evaluation metrics on the 2nd iteration.

Class	P	R	F1
Disciplines	0.93	1.0	0.96
Systems	1.0	0.97	0.98
Objects	1.0	0.92	0.96

present in the dataset but not extracted by the model. Since a gold standard was not available, the evaluation was performed manually by reviewing and correcting the model’s output.

Table 3: Information extraction task evaluation metrics.

Class	P	R	F1
Extracted entities	0.89	0.83	0.86

Conclusion and future development

This research presents a methodology aimed at automating cost estimation processes by leveraging LLMs to classify work items and extract relevant information from tender documents. The study specifically focuses on developing a dynamic prompting system framework to improve the accuracy of LLMs in this context. The methodology not only facilitates the creation of structured databases, it also aids in streamlining practitioners’ activities by automating a critical step of cost estimation.

One of the key contributions of this work is the integration of a cost-domain taxonomy that defines distinctions between classes within a hierarchical data structure. This taxonomy is fundamental not only for improving prompt performance, but also for supporting an iterative refinement of domain knowledge definitions. This iterative approach is crucial in cost management, as it helps contextualize the task within the specific domain. By continuously updating the dataset and adjusting the prompts accordingly, the methodology improves the performance of large language models, making them more effective in extracting and classifying information related to cost items. This approach improves the model performance and ensure that the models can accurately classify and extract details, even when confronted with minimum differences in object types or disciplines described in the work items.

The model was tested on a practical case study consisting of 70 work item descriptions, equally distributed across various object-related work items. Its performance was evaluated using the F1 metric, achieving a minimum score of 0.96. This value is expected to improve further with additional iterations and dataset refinement. Regarding the information extraction task, the model demonstrated lower accuracy, maintaining an F1 score of 0.86.

However, while these initial results are promising, further validation is needed. Future work should focus on test-

ing the methodology on larger datasets, including a wider variety of work item descriptions coming from different price lists documents, to ensure its scalability and robustness. Additionally, applying the methodology to different contexts, such as regulation checking, could provide valuable insights into its adaptability and performance across other sectors of the construction industry. Furthermore, future research involves the fine-tuning of the model to assess whether it is possible to achieve accurate results with fewer prompt steps, potentially reducing the complexity of the model, cost, time, and hence more accessible for real-world applications.

Data availability

The data presented in this study are openly available through GitHub under the MIT License.

The dataset is based on the 2023 edition of the Lombardy Regional Price List. As a result, both the dataset and the prompts are in Italian language.

Flat reinforced concrete and clay block slab with parallel ribs, cast in place with clay monoblocks. C20/25 concrete for ribs and screed, cast in place and vibrated. The upper slab with a thickness of no less than 5 cm, formwork, and temporary support scaffolding up to 4.50 m in height from the support plane, as well as the subsequent removal of the formwork. Excluded rebar reinforcement. - Total height: 17 cm (12 cm clay blocks + 5 cm slab).

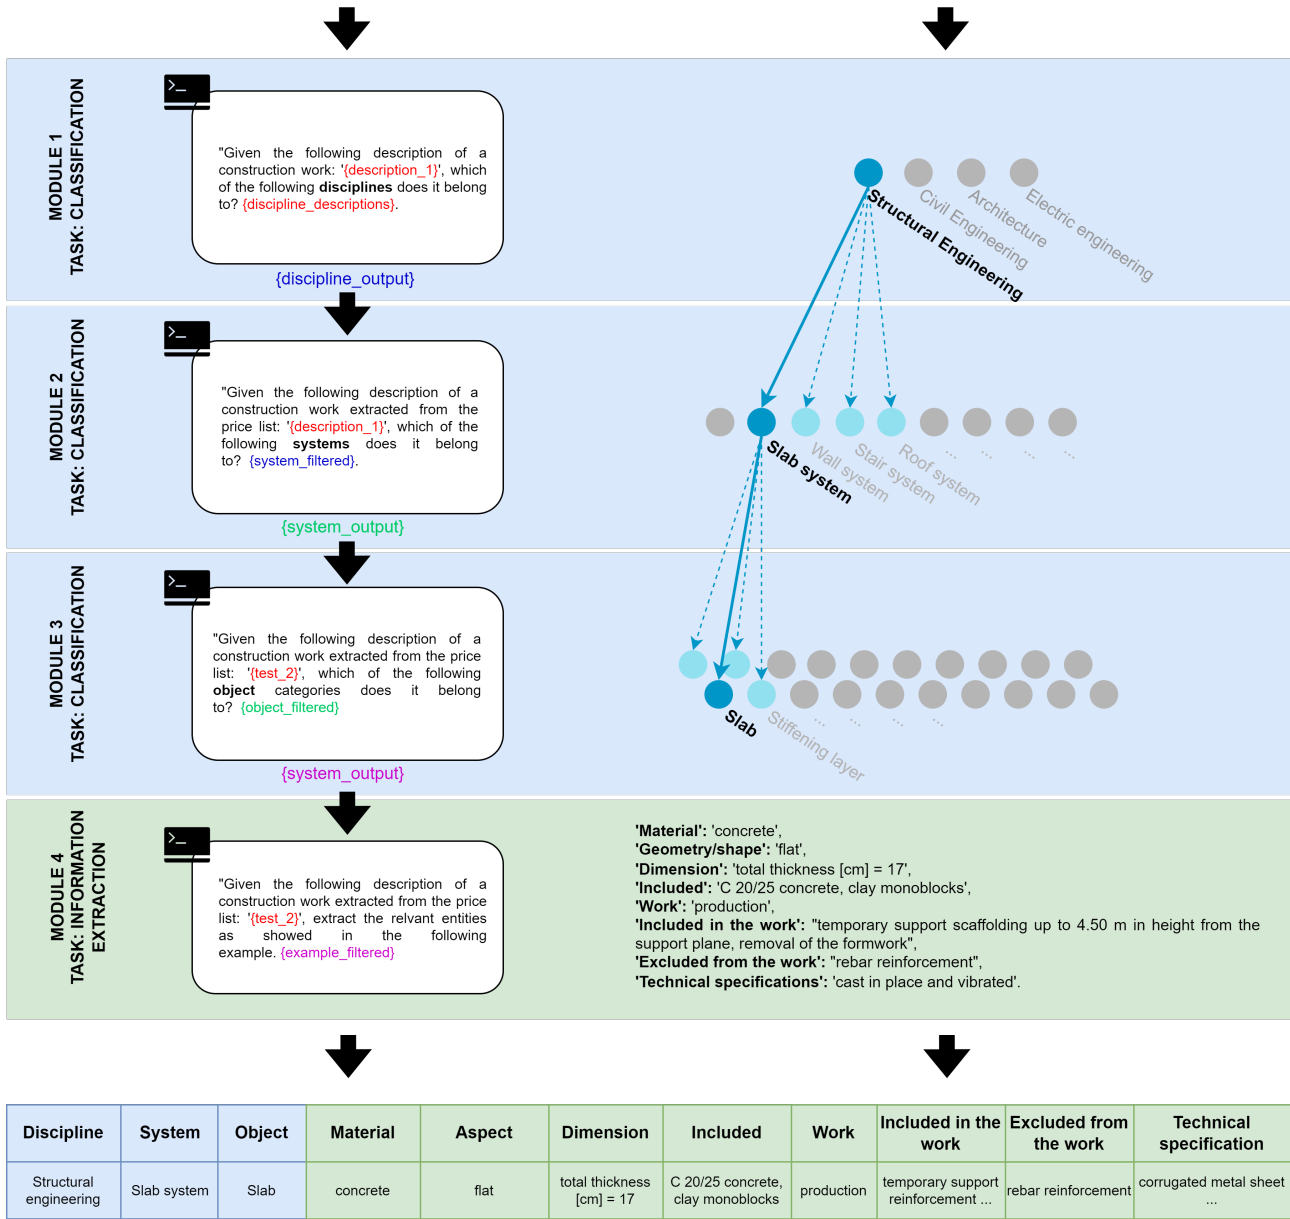


Figure 4: Practical example.

References

- Ashkenazi, O., Isaac, S., Giretti, A., Carbonari, A., and Durmus, D. (2023). Transforming building industry knowledge management: A study on the role of large language models in fire safety planning. pages 729–738.
- Cassandro, J., Donatiello, M. G., Mirarchi, C., Zanchetta, C., and Pavan, A. (2023). Reliability of ifc classes in ontology definition and cost estimation of public procurement. Proceedings of the 2023 European Conference on Computing in Construction and the 40th International CIB W78 Conference, 4:0–0.

- Cassandro, J., Gatto, C., Farina, A., Mirarchi, C., and Pavan, A. (2024). Innovative standardized cost data structure: application on price list document for estimating public tendering. Proceedings of the International Symposium on Automation and Robotics in Construction, pages 396–403.
- Chen, N., Lin, X., Jiang, H., and An, Y. (2024). Automated building information modeling compliance check through a large language model combined with deep learning and ontology. Buildings 2024, Vol. 14, Page 1983, 14:1983.

- Fuchs, S., Witbrock, M., Dimyadi, J., and Amor, R. (2024). Using large language models for the interpretation of building regulations.
- Gatto, C., Farina, A., Mirarchi, C., and Pavan, A. (2023). Development of a framework for processing unstructured text dataset through nlp in cost estimation aec sector. 4:0–0.
- Giuda, G. M. D., Giana, P. E., and Pattini, G. (2020). The shortening and the automation of payments: The potentiality of smart contract in the aeco sector. *Proceedings of International Structural Engineering and Construction*, 7:CON–12–1–CON–12–6.
- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220.
- Lee, J., Yi, J.-S., and Son, J. (2019). Development of automatic-extraction model of poisonous clauses in international construction contracts using rule-based nlp. *Journal of Computing in Civil Engineering*, 33:04019003.
- Martínez-Rojas, M., Marín, N., and Vila, M. A. (2013). A preliminary approach to classify work descriptions in construction projects. *Proceedings of the 2013 Joint IFSA World Congress and NAFIPS Annual Meeting, IFSA/NAFIPS 2013*, pages 1090–1095.
- Martínez-Rojas, M., Soto-Hidalgo, J. M., Marín, N., and Vila, M. A. (2018). Using classification techniques for assigning work descriptions to task groups on the basis of construction vocabulary. *Computer-Aided Civil and Infrastructure Engineering*, 33:966–981.
- Moon, S., Lee, G., Chi, S., and Oh, H. (2021). Automated construction specification review with named entity recognition using natural language processing. *Journal of Construction Engineering and Management*, 147.
- Pupeikis, D., Navickas, A. A., Klumbyte, E., and Seduikyte, L. (2022). Comparative study of construction information classification systems: Cci versus uniclass 2015. *Buildings* 2022, Vol. 12, Page 656, 12:656.
- Sepasgozar, S. M., Costin, A. M., Karimi, R., Shirowzhan, S., Abbasian, E., and Li, J. (2022). Bim and digital tools for state-of-the-art construction cost management. *Buildings* 2022, Vol. 12, Page 396, 12:396.
- Stas, S. and Abrishami, S. (2024). Decentralised automated bim collaboration: a blockchain and wbs integrated platform. *Smart and Sustainable Built Environment*, ahead-of-print.
- Tang, S., Liu, H., Almatared, M., Abudayyeh, O., Lei, Z., and Fong, A. (2022). Towards automated construction quantity take-off: An integrated approach to information extraction from work descriptions. *Buildings*, 12.
- Taye, M. M. (2010). Understanding semantic web and ontologies: Theory and applications. *Journal of Computing*, 2.
- Wan, H., Zhang, J., Chen, Y., Xu, W., and Feng, F. (2024). Generative ai application for building industry.
- Wang, J., Gao, X., Zhou, X., and Xie, Q. (2021). Multi-scale information retrieval for bim using hierarchical structure modelling and natural language processing. *J. Inf. Technol. Constr.*, 26:409–426.
- Wu, C., Li, X., Guo, Y., Wang, J., Ren, Z., Wang, M., and Yang, Z. (2022). Natural language processing for smart construction: Current status and future directions. *Automation in Construction*, 134:104059.
- Zheng, J. and Fischer, M. (2023). Dynamic prompt-based virtual assistant framework for bim information search. *Automation in Construction*, 155:105067.