Structuring Information in Government Documents Using Model-Driven Zero-Shot LLM Prompting

Stylianos Bourmpoulias[1,*], Dimitris Zeginis[1], Christos-Fanourios Patsouras[1] and Konstantinos Tarabanis[1]

[1]University of Macedonia, 156 Egnatia Street, 546 36, Thessaloniki, Greece

Abstract

Public administrations around the world produce large volumes of documents in many areas, including Human Resource Management (HRM). These documents include valuable data about people, positions, events, processes, and locations. However, the unstructured format of the documents makes it difficult to extract necessary information and limits their effective use by government agencies. The aim of the paper is to develop a semantic-driven HRM data model for education to enable a model-driven zero-shot Large Language Model (LLM) prompting approach for information extraction. Using various LLMs (e.g., Gemini, Llama, Claude4, ChatGPT), the data model is applied to administrative documents issued by the Greek Ministry of Education and published in DIAVGEIA.gov.gr, the national open government portal of Greece. Although at an exploratory stage, the study presents promising results regarding the capability of modeldriven LLMs in knowledge engineering, especially in text-intensive domains such as public administration.

Keywords

Public administration, Information extraction, Data model, HRM, LLMs, Ontology, Knowledge Graphs

1. Introduction

Public administrations around the world generate vast amounts of documents (e.g., laws, administrative decisions) across various functional areas (e.g., HRM). These documents include valuable data about people (e.g., education personnel), things (e.g., job positions), events (e.g., position assignments), processes (e.g., employment applications), and locations (e.g., school facilities). However, because this data is frequently provided unstructured, extracting the necessary information can be challenging. Therefore, public administrations are unable to properly utilize the potential of this important data.

The literature has introduced Generative AI, particularly LLMs, as effective tools for information extraction tasks [1]. Researchers have also investigated the integration of ontologies and Knowledge Graphs (KGs) using few-shot or even zero-shot learning techniques, showing that these can enhance LLM performance even without prior fine-tuning [2]. Ontologies and KGs assist LLMs in organizing recognized entities in a semantically coherent way by defining concepts and relationships within a particular domain and encoding them using RDF (Resource Definition Framework) triples. This can lead to the creation of a queryable semantic repository, i.e., a KG, that contains valuable information.

Although ontologies and KGs are frequently employed in education for pedagogical reasons (such as modeling learning domains), little is known about how they may be applied to education administration tasks, especially when managing educational staff [3]. Developing ontologies to assist LLMs in extracting valuable information from administrative documents (e.g., annual averages of teacher absences per course and school due to sick leave), could provide significant benefits for government agencies (e.g.,

PoEM2025: Companion Proceedings of the 18th IFIP Working Conference on the Practice of Enterprise Modeling: PoEM Forum, Doctoral Consortium, Business Case and Tool Forum, Workshops, December 3-5, 2025, Geneva, Switzerland

*Corresponding author.
†These authors contributed equally.
steliosbour@uom.edu.gr (S. Bourmpoulias); zeginis@uom.edu.gr (D. Zeginis); iis21046@uom.edu.gr (C. Patsouras); kat@uom.edu.gr (K. Tarabanis)
0009-0006-5288-3580 (S. Bourmpoulias); 0000-0002-5426-4031 (D. Zeginis); 0009-0008-7620-601X (C. Patsouras); 0000-0002-4663-2113 (K. Tarabanis)

optimized staffing) and ultimately enhance the quality of educational services offered to citizens (e.g., better teacher allocation).

Given the limited number of modeling initiatives in the HR domain of education, this study aims to develop a semantic-driven data model that captures the key elements of "Employment" and "Position Assignment" of education personnel suited to the Greek education system. This model will form the basis for applying a model-driven zero-shot LLM prompting approach to administrative documents issued by the Greek Ministry of Education agencies and published in DIAVGEIA.gov.gr, the national open government portal of Greece. This study is guided by the following research questions:

1. What are the fundamental building blocks of data that describe the employment and the position assignments of education personnel?
2. Can these building blocks be aligned with generic HRM data models, making the resulting HRM data model applicable to educational systems worldwide?
3. How effectively can Greek-capable LLMs (e.g., ChatGPT) structure unstructured data from Greek language administrative documents using the proposed HRM data model and a model-driven zero-shot LLM prompting approach?

The rest of the paper is structured as follows: Section 2 develops the background and related work of the study. Section 3 analyses the scope, goal, requirements, and procedure of developing the data model. Section 4 presents the experimental evaluation results and analysis. Finally, section 5 presents conclusions, study limitations, and recommendations for future work.

2. Background and related work

2.1. HRM data models and ontologies

The literature has proposed data models, that represent the key entities and relationships in the HRM domain. In particular, Hay [4] and Silverston [5] developed reference models offering generic patterns for modeling HR concepts such as "Employment" and "Position Assignment." These universal models provide standardized frameworks that organizations can adapt to their specific needs. Both authors view "Employee" not as an entity itself, but as a "role" assumed by a "Person" when employed by an organization. In contrast, Strohmeier and Röhr [6] introduce the "Employee" as a distinct entity with specific properties—thereby diverging from Hay's and Silverston's rolebased interpretation.

M. Jarrar et al. [7] called for applying semantic technologies in HRM and emphasized the need for domain-specific ontologies to enable knowledge-based automation (e.g., job search engines). The European Commission also developed an ontology that comprises thirteen modular ontologies as controlled vocabularies to describe job postings and job seekers' CVs [8] and introduced several controlled vocabularies for core HR concepts [9]. Additionally, the ESCO ontology was created to standardize the classification of skills, competencies, qualifications, and occupations to facilitate interoperability across EU countries and bridge the gap between education and the labor market [10].

HRM-specific data models and ontologies have also been developed in education. The European Commission has created controlled vocabularies (e.g., taxonomies, thesauri) for education-related concepts, including teaching staff [11]. In the U.S., the Common Education Data Standards (CEDS) initiative has established a Domain Entity Schema with hierarchies of domains, entities, categories, and elements covering the full education system, including HR areas such as K–12 staff employment and assignment [12]. While many ontologies exist in the education sector [3], few explicitly focus on HRM-related activities. Zemmouchi-Ghomari and Ghomari [13] and Rahman and Rabby [14] developed data specifications for Higher Education (HE), that define key HR concepts. Alrehaili et al. [15] introduced an HE ontology that integrates educational data into a structured, machine-readable format to automate tasks like academic staff–course allocation.

2.2. HRM Knowledge Graphs

According to the literature, KGs can integrate large volumes of data in a meaningful manner. To solve various issues in hiring, training, payroll, and HR systems, researchers in HRM have used KG technologies. In particular, HR-focused KGs have been created to assist with tasks like job market analysis [16], talent matching [17], recruitment automation [18], onboarding processes [19], and skills gap analysis [20]. While KGs have been established in education to support teaching and learning activities (e.g., curriculum design, personalized learning), few initiatives have concentrated on activities linked to education management. More precisely, Wang and Lin [21] designed a KG to address teaching arrangement challenges for cross-disciplinary professional courses. Similarly, I. Aliyu and S. Aliyu [22] proposed a KG to automate course-to-lecturer allocation in HE institutions. Lastly, Bourmpoulias et al. [23] introduce an entity-event KG for HRM in the public sector to reconstruct the evolution of education personnel in context and time.

2.3. Model-driven structuring of information using LLMs

According to the literature, there is a strong interplay between LLMs and models/ontologies/KGs [24], especially in text-intensive domains such as the public sector. For example in [2] authors propose the use of LLMs and ontologies as effective tools for information extraction tasks. More recently, GraphRAG [25] has been proposed as a novel RAG (Retrieval-Augmented Generation) approach that uses graphs (concepts/nodes, relationships/edges) as context for LLMs. GraphRAG, as a first step, uses LLMs to construct the graph, however the underlying data model for the graph construction is very simple (a flat list of concept/properties). Thus, resulting in not high quality or uniform results. On the contrary, when guided by model-driven prompts, LLMs are capable of consistently transforming unstructured text into structured data [26] even through zero-shot learning [2], [26]. The collaboration between LLMs, ontologies, and KGs enables KG construction, where LLMs generate candidate data and ontologies validate and organize it into triples. Approaches like [27] illustrate how LLMs can improve the accuracy of the extracted information.

2.4. The DIAVGEIA platform

In 2010, the Greek government established the national open government portal, "DIAVGEIA" (https://diavgeia.gov.gr/) (meaning "clarity"), to promote government accountability. Since then, all government institutions have been required to upload their acts and decisions (e.g., budgets, appointment decisions) in the platform to be accessible to the public. Each uploaded document is assigned a unique ID number to certify its authenticity and legality. As of now, 68.4 million documents have been uploaded to DIAVGEIA, with an ongoing rate of 16,000 decisions being published each working day. Specifically, in 2024, agencies from the Greek Ministry of Education uploaded 54,305 administrative decisions including HRM-related acts.

3. An HRM Data Model for Education

3.1. Scope and goal

The study scopes the administrative acts that are published in DIAVGEIA and pertain to the employment of education personnel and their assignment to primary and secondary public education positions. It encompasses all professionals directly involved in the educational process (e.g., teachers, psychologists, social workers) through all types of employment (e.g., regular, substitute). Positions like school headmasters, which include administrative and instructional duties, are also included. The study aims to create a semantic-driven data model that identifies the fundamental building blocks of data (such as actors, events, and locations) related to educational staff's employment and position assignments. The classes and their properties will give LLMs the metadata they need to organize the textual data from DIAVGEIA documents meaningfully. This way, the model will help with information extraction using a model-driven zero-shot LLM prompting strategy, resulting in better model-driven prompts and more organized results.

3.2. Requirements

To ensure alignment with the study's overall goals, we established a set of modeling requirements. The four fundamental principles outlined by [28] were used to engineer the modeling requirements systematically. First, the model should be accurate for the purpose at hand, thus guiding LLMs to automatically extract and structure textual information from DIAVGEIA (Req. 1 – Validity). Second, it should be aligned with generic HRM data models and, at the same time, represent the specifics of the Greek public education system. In particular, in Greece, teachers are classified according to their qualifications into specific branches (e.g., mathematics), levels (e.g., primary/secondary education), and education types (e.g., general/special education). The Ministry of Education centrally hires them and assigns them to particular educational regions within the country. Depending on local educational needs, the regional agencies (e.g., Directorates of Primary Education) assign them to one or more positions (Req. 2 – Credibility). Third, it should enable the transformation of the unstructured textual data into structured, queryable datasets (Req. 3 – Utility). Fourth, it should be implementable within time, resources, and data constraints, relying on the lightweight and scalable use of pre-trained language models, without domain-specific finetuning (Req. 4 – Feasibility).

3.3. Modeling procedure

Our modeling procedure is grounded in modeling hierarchy theory [29], which defines a four-layer meta-model architecture. The raw data at Level 0 serves as the foundation for the Level 1 model. At Level 2, the model is abstracted into a meta-model that organizes and classifies the Level 1 components. At the top, Level 3 comprises the meta^2-model, which defines the foundational concepts used to develop the lower levels. Moving from Level 0 to Level 3 involves a process of classification, while moving downward follows a process of instantiation.

The modeling hierarchy schema provides a solid framework and the necessary mechanisms to align our modeling procedure with the previously outlined requirements. To this purpose, we adopted a combination of bottom-up and top-down steps (see Fig. 1). Rather than following a linear sequence, these two approaches are applied in a complementary manner. While the top-down approach helps us focus on high-level system needs, the bottom-up approach allows us to construct models from lower-level components.

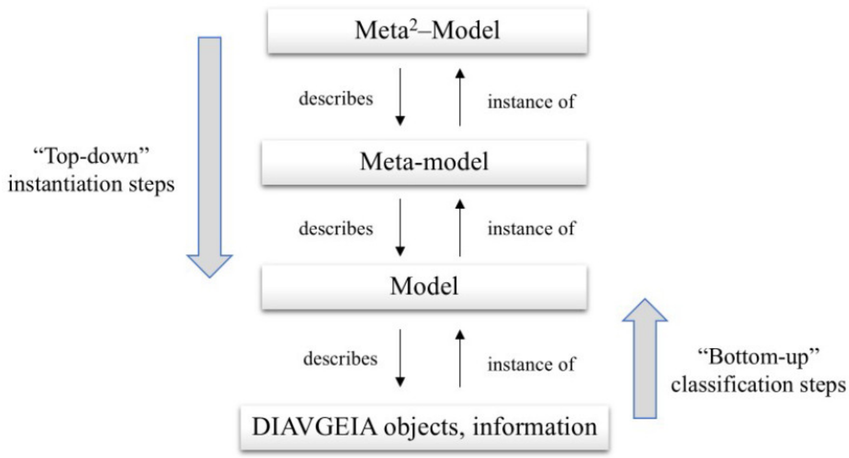

Figure 1: Modeling procedure for developing the HRM Data Model, inspired by D. Karagiannis and H. Kühn [29]

More precisely, we created the model according to the recommendations made by [30]. After establishing the domain and scope, we chose generic HRM data models from the literature examined in Section 2. We also collected documents from DIAVGEIA and their underlying legal texts related to employment and assignments of education personnel. Our review, supported by our experience as domain experts, enabled a thorough understanding of the domain. After collecting key terms as the basis of our ontology, we selected terms with independent existence and organized them into a class hierarchy. We then analyzed the remaining terms to define class properties (e.g., value types, allowed values, cardinality). Throughout this process, we aimed to align our ontological building blocks with generic HRM data models. Finally, we created individual instances of each class.

Using a bottom-up approach to describe data from DIAVGEIA documents was considered suitable, as it can enrich LLMs with detailed metadata for organizing unstructured text (Req. 1). This also supports prompt optimization through the application of ontology-based prompts, which help LLMs extract key information and convert it into triples (Req. 3). This process can lead to a queryable KG supporting advanced analytics (Req. 3). Moreover, drawing from generic HR data models and refining them into domain-specific models minimizes bias and enhances overall model reliability (Req. 2).

3.4. Design and development

Using our methodology, we initially searched DIAVGEIA for documents issued by agencies of the Greek Ministry of Education, scoping the employment and assignment of educational personnel. Due to the inconsistent terminology used across government agencies [31], we also reviewed the overarching legal texts. Afterward, we created an initial list of terms such as "teacher", "special education", and "position". We then classified them according to the six Zachman interrogatives [32]. This process resulted in the development of an initial controlled vocabulary that encompasses key business elements (e.g., position), processes (e.g., applying for an appointment), locations (e.g., region of appointment), educational organizations (e.g., Directorate of Secondary Education), events (e.g., position assignment), and goals (e.g., meeting teaching needs).

We then identified terms representing entities with independent existence (e.g., "secondary education teacher") and terms that describe these entities (e.g., "substitute"). The former were classified as classes within the ontology, while the latter were considered as their attributes. Next, we organized the classes hierarchically. For instance, "Position type" was defined as a super-class of "Secondary education teacher". In contrast, more specific concepts (e.g., "Secondary education mathematics teacher"), were classified as sub-classes of broader categories and positioned at the bottom. Finally, we attached slots to the most general classes and specified their characteristics. This process produced the core building blocks for an HRM data model representing the employment and position assignments of educational personnel (see Fig. 2).

Special care was taken to ensure that the developed building blocks are aligned with established generic HRM data models proposed in the literature, such as those referenced in [4] and [5]. This alignment could enhance their adaptability across various educational systems worldwide. More specifically, two main actors were identified: “Person” and “Organization.” Since both share common attributes (e.g., name, address), a unified super entity called “Party” was established. Individuals and organizations take on specific roles: a person functions as an “Employee,” while an organization serves as an “Internal Organization”. Consequently, the “Employment” entity is modeled as a sub-type of “Party Relationship,” which represents the relationship between a specific person and an organization (e.g., Ministry of Education). This entity also includes a “from date”, i.e., the hire date and a “thru date”, i.e. the end of employment.

A person employed by an organization is assigned, for a specific period of time, to a particular position (e.g., mathematics teacher). In the Greek education system, teaching positions are not linked to the hiring organization (e.g., Ministry of Education) but usually to specific schools (e.g., 1st High School of Athens). Therefore, the “Position” entity is modeled as distinct from

“Employment” and is directly linked to the “Organization” through a “defined by” relationship [33], [4]. The assignment itself is represented by a “Position Assignment” entity, which includes various types of assignments (e.g., secondment) and attributes such as start and end dates.

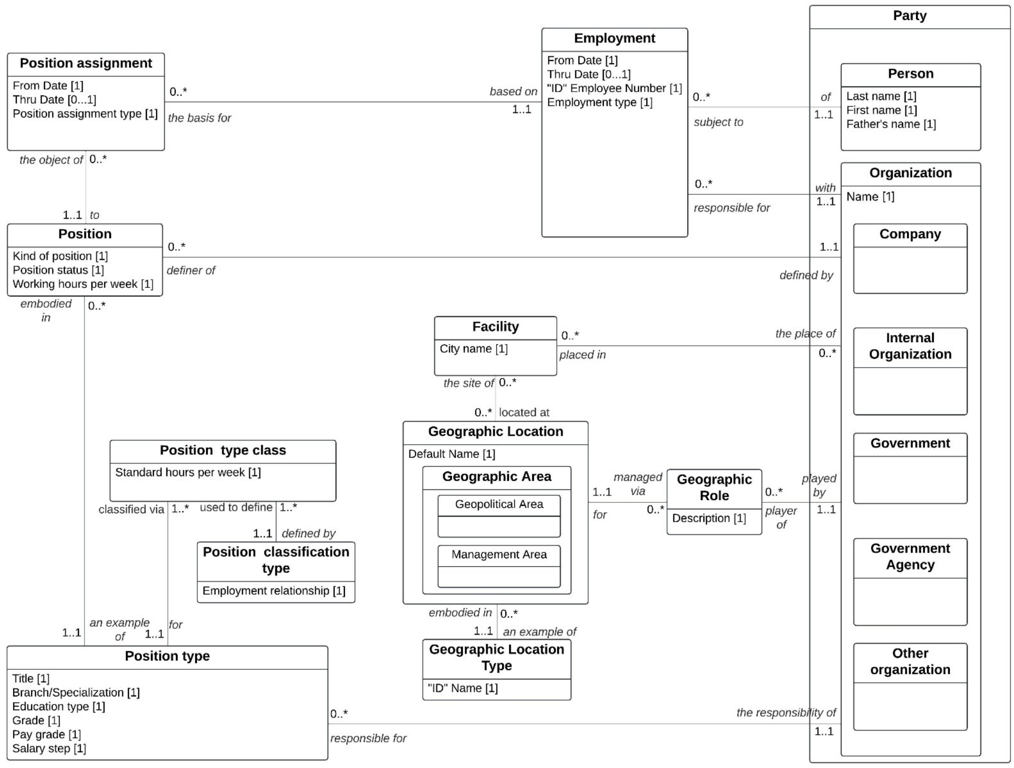

Figure 2. The HRM Data Model for Primary and Secondary Education based on [4], [5], [33].

Since many positions share common characteristics (e.g., job title), the entity “Position Type” is introduced [4], [5], [33]. Each position must belong to exactly one position type, which is managed by a single organization [4]. To facilitate further classification—like classifying a full-time substitute special education mathematics teacher—the model includes “Position Classification Type” proposed by [5]. This categorizes how the position will be compensated, such as whether it is paid hourly or offered on a temporary basis. Additionally, "Position Type Class" serves as an intersection between "Position Type" and "Position Classification Type," enabling more detailed groupings [5].

While the concept of "space" is not traditionally regarded as a building block of HR models in the literature, it plays a significant role in national-level education staffing. The Ministry of Education centrally manages recruitment, but teachers are assigned to regional organizations (e.g., Directorates of Secondary Education), which are in charge of specific management areas and assign teachers to positions. The Ministry of Education establishes the boundaries of these areas, which may not align with the nation's recognized geopolitical boundaries. Schools are also classified according to their geographic location. As a result, factors such as a teacher’s years of service in remote geographic areas (e.g., islands) receive special consideration in administrative processes, including transfer eligibility.

Because of the significant role of geography in education personnel management, we introduced the concept of “geographic location” into the HR model. Based on [4], we adopted the concepts of "Management area", a sub-type of geographic location defined by boundaries set by the organization, and "Facility" to represent the physical locations of schools. Additionally, we introduced the concept of "Geographic role" to represent the delegation of education personnel management to regional agencies of the Ministry of Education, each responsible for a defined management area.

4. Experimental results and analysis

This section presents the results of experimental evaluations on using LLMs to extract structured information from DIAVGEIA documents through a model-driven approach. The experiment aimed to assess how effectively Greek-capable LLMs can structure unstructured government data. The experiment was conducted in seven phases. *First*, we developed an initial prompt template to be used for information extraction from documents describing teachers' position assignments. The prompt is model-driven, meaning that the HRM data model structure is embedded in the prompt and requests results based on this specific format, to enforce semantic consistency in LLM outputs.

Second, we defined the test set of documents to apply the model-based prompt by searching DIAVGEIA for various types of position assignments (e.g., secondment) and employment relationships (e.g., regular, substitute) of teachers. Given that teachers can hold multiple positions simultaneously, we included a range of documents, including those detailing the assignment of several teachers to various roles. Given the variability in wording and terminology across public organizations, we selected documents from 14 public entities, comprising a total of 36 documents—18 involved assigning each teacher to a single position, and 18 involved assigning each teacher to multiple positions. All documents were public records, containing only basic information such as names and fathers' names, and adhered to GDPR requirements, excluding sensitive personal data.

Some interesting remarks about the documents (see Fig. 3) include: i) most of the documents refer to multiple (some times more than 50) position assignments, ii) the documents usually begins with a preamble text that refers to all assignments, including info about things like the start and the end date, iii) after the preamble text follows a table with specific information for each assignments such as information about the person, management area, working hours, iv) in some cases the table lists more than one assignment for the same person (e.g., she/he may be assigned at two schools on a part-time basis). These structural characteristics make information extraction particularly challenging.

Figure 3. Sample "Position assignment" document with relevant annotations (in Greek)

Third, we manually annotated the 36 documents based on the HRM data model. This manually generated information was regarded as the ground truth. *Fourth*, we applied the prompt format to LLMs that understand Greek, and we had access. Seven LLMs were employed: Claude 3.5, Claude 3.7, Claude 4, Llama (lama 3.3 30b parameters), Deepseek- R1, ChatGPT 4, and Gemini 2.5 Pro. The documents were first converted from PDF to text, and then the extracted table was converted to a dataframe to have more structured text. The LLMs produced output in triples (e.g., Person_1 has_first_name Chris), which were evaluated against the ground truth.

Fifth, we improved the prompt after identifying several extraction errors during the initial phase of the experiment across all LLMs. These errors included failing to identify the correct start and end dates and omitting multiple position assignments. The errors were due to the documents' complex structure, which contained multiple dates (such as the document issue date and the council opinion date), and to the domain-specific semantic complexity. To improve the accuracy of the returned information, we decided to enhance the prompt by adding a section with domain-specific logic and rules. This led to the development of the final prompt template used for evaluation (see Fig. 4).

```
< Separate the piece of <Greek Text> extracted from a position assignment decision exclusively into fields with the
following <Format>, taking into account the <Descriptions> and <Logic_and_Rules> of the fields. Return the result
exclusively using the provided <Format>. Do not enter fields for which you have not found values.>
<Format>
Person_{I} has_last_name <Last_Name>
Person_{I} has_first_name <First_Name>
Person_{I} has_father's_name <Father's_Name>
Person_{I} has_ID_employee_number <ID_Employee_Number>
Person_{I} is_subject_to Position_assignment_{K}
Position_assignment_{K} has_from_date <From_Date>
…
<Descriptions>:
Last name: <The one and only one last name of a person>
…
<Logic_and_Rules>:
…
Each Position must be an example of one and only one Position Type (e.g. "Εκπαιδευτικός").
…
```

Figure 4. Excerpt of the prompt template (complete version available at 10.5281/zenodo.17615179)

Sixth, we applied the final version of the prompt to the LLMs and, *seventh*, we compared the returned AI-generated triples against the ground truth. The evaluation was performed semi-automatically due to the large number of triples per document. An automatic script first normalized the values (e.g., by removing symbols and accents) and then identified identical triples. The script calculated a similarity score for triples with value differences. If the score was above 85%, the triple was considered correct. After the script returned the correct triples, we double-checked them manually to correct any remaining mistakes. Ground truth and AI-generated triplets that were the same were considered correct. Minor differences, such as special symbols (e.g., <>, ‘, ", etc.) or variations in wording/case of values (e.g., 'Western', 'western'), were not taken into account and were also considered correct.

Through the evaluation process, we identified: i) True Positives: Triples correctly identified by AI, ii) False Positives: Incorrectly added AI-triples, and iii) False Negatives: Missed triples by AI. Finally, we calculated Precision, Recall, and F1-score (check Table 1 and Fig. 5) using the following formulas:

- *Precision = True Positives / (True Positives + False Positives)*
- *Recall = True Positives / (True Positives + False Negatives)*
- *F-1 = 2*(Precision*Recall)/ (Precision + Recall)*

Table 1
Precision, Recall and F-1 metrics for the whole sample of documents

	Gemini	ChatGPT	Llama	Deepseek	Claude3.7	Claude3.5	Claude4
Precision	0,913	0,774	0,657	0,814	0,856	0,845	0,841
Recall	0,889	0,694	0,532	0,699	0,673	0,687	0,768
F-1	0,898	0,719	0,573	0,739	0,734	0,744	0,790

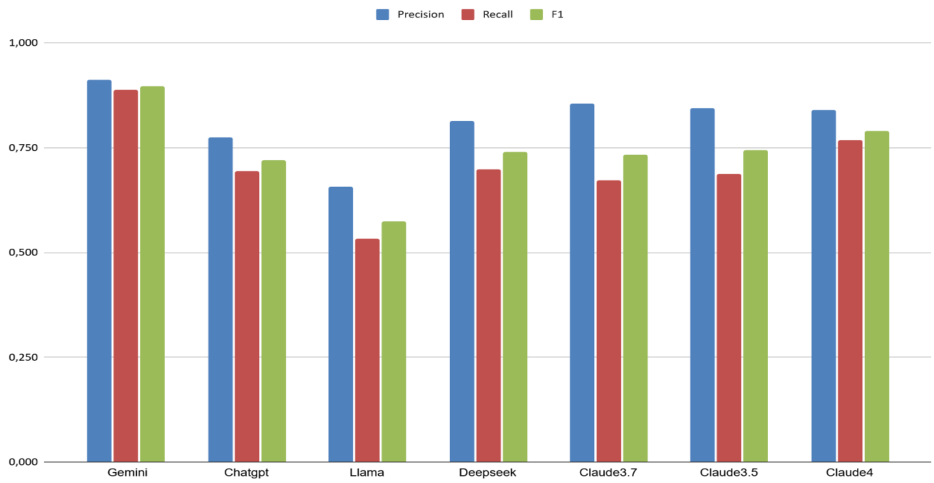

Figure 5. Precision, Recall and F-1 metrics

According to the results, Gemini achieved the highest Precision (0.913), with Claude's versions hovering around 0.841-0.856. Deepseek and ChatGPT also perform well with scores of 0.814 and 0.774, respectively. In contrast, with 0.657, Llama has the lowest Precision. Regarding Recall, Gemini has the highest score (0.889), while Claude4, with 0.768, shows a remarkable improvement over previous versions, around 0.673-0.687. Deepseek's recall (0.699) is similar to ChatGPT's (0.694) and better than versions 3.7 and 3.5 of Claude. With 0.532, Llama has the lowest recall, which means it misses a significant number of cases. Overall, among the LLMs tested, Gemini 2.5 Pro and Claude4 performed the best, with Gemini 2.5 Pro achieving the top F1-score (0.898). Table 2 and 3 present separate scores for documents describing “one-to-one” and “one-to-many” person-position assignments respectively. In general, the “one-to-many” documents are more complex than the “one-to-one” and this is reflected in the scores. Except for Gemini 2.5 Pro, all LLMs struggled to extract information from “one-to-many” documents, achieving lower scores.

Table 2
Precision, Recall, and F-1 for documents describing “one-to-one” person-position assignments

	Gemini	ChatGPT	Llama	Deepseek	Claude3.7	Claude3.5	Claude4
Precision	0,903	0,838	0,694	0,853	0,888	0,905	0,904
Recall	0,893	0,855	0,712	0,821	0,805	0,831	0,867
F-1	0,893	0,818	0,665	0,815	0,802	0,829	0,885

Table 3
Precision, Recall, and F-1 for documents describing “one-to-many” person-position assignments

	Gemini	ChatGPT	Llama	Deepseek	Claude3.7	Claude3.5	Claude4
Precision	0,921	0,716	0,623	0,780	0,828	0,792	0,784
Recall	0,885	0,552	0,372	0,590	0,556	0,559	0,680
F-1	0,899	0,608	0,459	0,654	0,641	0,639	0,706

Regarding the information extraction errors, serious issues were identified in cases where a teacher holds two or more positions. While Gemini and Claude4 perform well, other LLMs often omit position assignment information, which is reflected in the Recall values which are lower than Precision. Except for Gemini, LLMs also face challenges in accurately identifying position assignments' dates. The semantic complexity (e.g., simultaneous presence of multiple dates) and the conceptual ambiguity (e.g., lack of explicit reference to the start and end dates) of the documents often led LLMs to incorrect results. Finally, issues such as ambiguities in table parsing (e.g., unclear cell boundaries and structure, and header misinterpretation) made it even more difficult for LLMs to identify information correctly. Among these factors, table parsing ambiguity (e.g., simultaneous presence of multiple values in the table header) often leads to challenges in accurately identifying the correct position assignment type, the kind of position, and the working hours per week.

5. Conclusions and discussion

The primary objective of this study was to structure unstructured government data in the field of Human Resource Management (HRM) using a model-driven zero-shot LLM prompting approach. To this end, the study developed a semantic-driven HRM data model tailored to the specifics of public education. This model captures the fundamental building blocks of data associated with the concepts of “Employment” and “Position Assignment”. Beyond key HR building blocks—such as employment, position, position type, and position assignment—the model systematically integrated the spatial dimension, often overlooked in HRM ontologies, to address the complexities of national-level education staffing. Although the model was specifically designed for the context of public education in Greece, it is closely aligned with generic HRM data models, allowing for adaptability in various educational systems worldwide. Furthermore, the model was embedded in a prompt template that ensures semantic consistency in the outputs generated by the LLM.

The study used real-world, high-volume government data from thirty-six (36) documents issued by fourteen (14) regional government agencies of the Greek Ministry of Education and published on DIAVGEIA.gov.gr. These documents scope the assignment of sixty-four (64) teachers to various positions. More specifically, they contain valuable information about teachers’ personal information (e.g., first and last name), position assignments (e.g., start and end date), position types (e.g., job title, employment relationship), schools that define the positions, education management areas, and competent authorities having jurisdiction over them. To assess how effectively Greek-capable LLMs could extract and structure this information, the study utilized seven competitive Greek-capable LLMs for evaluation. The assessment results are promising since most LLMs achieved good scores. However, in more complex cases, there is still room for improvement.

A limitation of the study is that it focuses on the employment and the position assignment of teachers, which are the most frequently categories in the HR-centric documents of the Greek Ministry of Education published on DIAVGEIA. As future work, we plan to expand the model by incorporating additional, less frequently published concepts. These concepts cover various HRM-related processes, such as benefit payments and employee terminations. We also plan to use and evaluate other prompting approaches (few-shot, chain of thought, etc.) and prompt designs (e.g., provide the structure of the ontology in the prompt as RDF or JSON, Pydantic code). Lastly, we aim to leverage European open-access LLMs, such as the Mistral model, Aleph Alpha Luminous, and other models emerging from EU research initiatives.

Although still in its early stages, the research shows encouraging results regarding the interaction between LLMs, domain-specific models, and KGs. In particular, it demonstrates the effectiveness of a model-driven zero-shot LLM prompting approach in extracting structured information from a large amount of unstructured government data. This method can help public administration worldwide to unlock the full value of its dispersed, document-based data. Moreover, the demonstrated capability of model-driven LLMs in knowledge engineering suggests strong potential for broader impact, especially in text-intensive domains such as public administration.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly and ChatGPT in order to: Paraphrase and reword, Improve writing style and Grammar and spelling check. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] D. Xu et al., Large Language Models for Generative Information Extraction: A Survey, Frontiers of Computer Science 18.6 (2024). doi: 10.48550/arXiv.2312.17617.

[2] D. Zeginis, E. Kalampokis, and K. Tarabanis, Applying an ontology-aware zero-shot LLM prompting approach for information extraction in Greek: the case of DIAVGEIA gov gr, in: Proceedings of the 28th Pan-Hellenic Conference on Progress in Computing and Informatics, in PCI '24. New York, NY, USA: Association for Computing Machinery, pp. 324–330. doi: 10.1145/3716554.3716603.

[3] K. Stancin, P. Poscic, and D. Jaksic, Ontologies in education – state of the art, Educ Inf Technol, vol. 25, no. 6 (2020) 5301–5320. doi: 10.1007/s10639-020-10226-z.

[4] D. C. Hay, Enterprise Model Patterns: Describing the World (UML Version). Denville, NJ, USA: Technics Publications, LLC, 2011.

[5] L. Silverston, The Data Model Resource Book, Volume 1: A Library of Universal Data Models for All Enterprises. John Wiley & Sons, 2011.

[6] S. Strohmeier and F. Röhrs, Conceptual Modeling in Human Resource Management: A Design Research Approach, AIS Transactions on Human-Computer Interaction 9. 1 (2017) 34–58.

[7] M. Jarrar, L. Vervenne & D. Maynard, HR-Semantics Roadmap - The Semantic Challenges and Opportunities in the Human Resources Domain, Technical Report, 2007.

[8] Ontology Engineering Group (OEG), Human Resources Management Ontology, 2013. URL: https://interoperable-europe.ec.europa.eu/collection/eu-semantic-interoperabilitycatalogue/solution/human-resources-management-ontology.

[9] European Union, EU Vocabularies. URL: https://op.europa.eu/en/web/euvocabularies/concept/-/resource?uri=http://eurovoc.europa.eu/100153.

[10] Johan De Smedt (TenForce). The ESCO ontology. Revision: 2.0.0. URL: http://data.europa.eu/esco/model/2.0.0

[11] European Union, EU Vocabularies. URL: https://op.europa.eu/en/web/euvocabularies/concept/-/resource?uri=http://data.europa.eu/bkc/006.06&lang=en

[12] The CEDS Initiative, Common Education Data Standards (CEDS). URL: https://ceds.ed.gov/Default.aspx

[13] L. Zemmouchi-Ghomari and A. Ghomari, Towards a Reference Ontology for Higher Education Knowledge Domain, International Review on Computers and Software 8.2 (2013) 474-488.

[14] M. Rahman and G. Rabby, Design and Development of a University Human Resource Ontology Model for Semantic Web, International Journal of Computer Science and Network Security 17 (2017) 187–182.

[15] N. A. Alrehaili et al., Ontology-Based Smart System to Automate Higher Education Activities, Complexity 2021.1 (2021). doi: 10.1155/2021/5588381.

[16] R. Boselli, M. Cesarini, F. Mercorio, and M. Mezzanzanica, Classifying online Job Advertisements through Machine Learning, Future Generation Computer Systems 86 (2018) 319–328. doi: 10.1016/j.future.2018.03.035.

[17] D. Liu, Y. Yang, M. Zhang, W. Wu, X. Xie, and G. Sun, Knowledge Enhanced Multi-Interest Network for the Generation of Recommendation Candidates, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, in: CIKM '22. New York,

NY, USA: Association for Computing Machinery, Oct. 2022, pp. 3322–3331. doi: 10.1145/3511808.3557114.
[18] U. Shahbaz, A. Beheshti, S. Nobari, Q. Qu, H.-Y. Paik, and M. Mahdavi, iRecruit: Towards Automating the Recruitment Process, in Service Research and Innovation, H.-P. Lam and S. Mistry, Eds., in Lecture Notes in Business Information Processing. Cham: Springer International Publishing, 2019, pp. 139–152. doi: 10.1007/978-3-030-32242-7_11.
[19] E. Ritz, F. Donisi, E. Elshan, and R. Rietsche, Artificial Socialization? How Artificial Intelligence Applications Can Shape A New Era of Employee Onboarding Practices: 56th Annual Hawaii International Conference on System Sciences, HICSS 2023," Proceedings of the 56th Annual Hawaii International Conference on System Sciences, HICSS 2023, pp. 155–164, 2023, doi: 10.24251/HICSS.2023.020.
[20] C. H. Loh, S. Wei, C. T. Phua, A. Mohd, A. Tan, and B. K. Seow, Nacelle: Knowledge Graph-Based Conversational AI for Skills Gap Analysis to Achieve Sustainable Learning at Workplace" in: Proceedings of 2nd International Conference on Artificial Intelligence: Advances and Applications: ICAIAA 2021. Singapore: Springer Nature Singapore, 2022. p. 359-371. doi: 10.1007/978-981-16-6332-1_31.
[21] P. Wang and J. Lin, Construction and Application of Knowledge Graph of Teaching Arrangement for Cross-disciplinary Majors, in: 2021 IEEE 3rd International Conference on Computer Science and Educational Informatization (CSEI), June 2021, pp. 85–88. doi: 10.1109/CSEI51395.2021.9477719.
[22] I. Aliyu, K. A. F. D., and S. Aliyu, Development of Knowledge Graph for University Courses Management, IJEME 10.2 (2020) 1-10. doi: 10.5815/ijeme.2020.02.01.
[23] S. Bourmpoulias, D. Zeginis, and K. Tarabanis, An Entity Event Knowledge Graph for Human Resources Management in Public Administration: the case of education personnel, in: 2023 IEEE 25th Conference on Business Informatics (CBI), June 2023, pp. 1–8. doi: 10.1109/CBI58679.2023.10187510.
[24] R. Buchmann et al., Large language models: Expectations for semantics-driven systems engineering" Data & Knowledge Engineering 152 (2024). doi: 10.1016/j.datak.2024.102324.
[25] D. Edge et al., From Local to Global: A Graph RAG Approach to Query-Focused Summarization, arXiv preprint arXiv:2404.16130 (2024). doi: 10.48550/ARXIV.2404.16130.
[26] Z. van Cauter and N. Yakovets, Ontology-guided Knowledge Graph Construction from Maintenance Short Texts, in: Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models (KaLLM 2024). 2024. p. 75-84. doi: 10.18653/v1/2024.kallm-1.8.
[27] J. H. Caufield et al., Structured Prompt Interrogation and Recursive Extraction of Semantics (SPIRES): a method for populating knowledge bases using zero-shot learning, Bioinformatics 40.3 (2024). doi: 10.1093/bioinformatics/btae104.
[28] S. Robinson, Conceptual modelling for simulation Part I: Definition and requirements, Journal of the Operational Research Society 59 (2008) 278–290. doi: 10.1057/palgrave.jors.2602368.
[29] D. Karagiannis and H. Kühn, Metamodelling Platforms, In: EC-web. 2002. p. 182. doi: 10.1007/3540-45705-4_19.
[30] N. F. Noy and D. L. McGuinness, Ontology Development 101: A Guide to Creating Your First Ontology, p. 25, 2001.
[31] S. Bourmpoulias, D. Zeginis, and K. Tarabanis, Towards a Common Reference Infrastructure for Documenting Public Sector Administrative Processes: The Case of 'MITOS', in: Companion Proceedings of the 17th IFIP WG 8.1 Working Conference on the Practice of Enterprise Modeling Forum, M4S, FACETE, AEM, Tools and Demos co-located with PoEM 2024, Stockholm, Sweden, December 3-5, 2024.
[32] Zachman, Enterprise Architecture, The Issue of the Century. Database programming and design 10.3 (1997) 44-53.
[33] D. C. Hay, Data Model Patterns: Conventions of Thought. Dorset House Pub., 1996.