

LLM-Assisted Ontology Engineering and Construction of a French Legal Knowledge Graph

Génesis Montenegro^{1,2}, Mokhtar Boumedyen Billami², Catherine Faron¹, Fabien Gandon¹ and Pierre Monnin¹

¹Université Côte d'Azur, Inria, CNRS, I3S, Sophia-Antipolis, France

²Berger-Levrault, Labège, France

Abstract

Maintenance regulations are complex legal texts that are difficult to exploit when addressing a specific case and challenging to integrate into operational systems. This paper presents a two-stage LLM-assisted workflow for French maintenance regulations: ontology engineering from a SEMLEG-based core ontology, followed by construction of an ontology-grounded French legal knowledge graph. The first stage consists in the open extraction of typed entities and triples from a stratified corpus sample, the normalization of labels through embedding-based fusion, and the induction of candidate object properties with their signature (domain and range). The second stage uses the resulting ontology to guide the closed extraction of triples and RDF graph construction over the full corpus. Experiments with GPT-4 . 1 and mistral-large-2512 show robust structured outputs, near-complete class alignment, and a substantial reduction of duplicated entities and predicates after fusion. Fewer than 20% of triples introduce unseen properties, while lower exact signature compliance reveals new domain-range combinations for existing predicates. These results point to predicate normalization and the validation of newly observed relation signatures as key refinement steps for industrial maintenance settings.

Keywords

knowledge graph construction, ontology engineering, language models, legal texts, maintenance regulation

1. Introduction

Regulatory compliance in industrial maintenance requires the interpretation of complex and evolving legal texts, making it difficult to exploit them in operational systems such as Computerized Maintenance Management Systems (CMMS). Transforming such texts into structured knowledge is therefore a key step toward queryable, reusable, and interoperable legal information. Recent work in legal information extraction and relation extraction has shown the potential of large language models (LLMs) for identifying entities and relations in legal documents [1, 2]. However, most existing approaches remain dataset-driven and do not provide an end-to-end pipeline for moving from legal text to an ontology-grounded knowledge graph. Ontology-based resources offer useful foundations for structuring legal and maintenance knowledge. The European Legislation Identifier (ELI) supports interoperable descriptions of legal resources, while SEMLEG models regulatory rules [3, 4]. In the maintenance domain, existing ontologies capture activities, assets, procedures, and industrial processes [5, 6, 7]. Nevertheless, these resources do not directly provide the domain-specific relation vocabulary required to populate a legal knowledge graph for maintenance regulations.

To address this gap, we present an end-to-end LLM-assisted pipeline for moving from texts describing French maintenance regulations to an ontology-grounded legal knowledge graph. The workflow comprises two stages: (1) ontology engineering, where a SEMLEG-based core ontology is enriched through LLM-based extraction, embedding-based normalization, and signature-driven property induction; and (2) knowledge graph construction, where the resulting ontology guides triples extraction from the full French regulatory corpus. We contribute a compact workflow combining LLM flexibility with ontology-based structural constraints, together with a preliminary evaluation of the resulting ontology and KG.

Posters, Demos, Blue Sky, and Tutorials at SEMANTiCS 2026, Sep 2026, Ghent, Belgium

0009-0002-9354-2265 (G. Montenegro); 0000-0003-4428-4298 (M. B. Billami); 0000-0001-5959-5561 (C. Faron); 0000-0003-0543-1232 (F. Gandon); 0000-0002-2017-8426 (P. Monnin)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. LLM-Assisted Ontology Engineering

Semantic Scope and Corpus. The semantic scope is defined through competency questions targeting maintenance obligations: involved actors, affected artifacts, contextual conditions, and legal sources. We start from SEMLEG and retain the classes required by the extraction workflow, including *Actor*, *Action*, *Artifact*, *Condition*, *Source*, *Location*, *Reason*, *Situation*, and *Time*, together with the object properties whose domain and range correspond to these classes. The corpus is built from official French legal texts retrieved from Légifrance¹ and filtered using maintenance-oriented references from the Apave regulatory guide.² These references are parsed and normalized before retrieving the corresponding legal articles and metadata, yielding a focused corpus of 6,370 regulatory articles covering 20 topics. The two stages of the workflow use this corpus differently: ontology engineering is performed on a stratified sample of 1,389 articles, corresponding to approximately 10% of each domain–document title pair, while KG construction (ontology population) is performed over the complete corpus.

Open Class-Guided Triple Extraction. Each sampled article a_i is processed with three prompts. A first prompt π_i^{ent} extracts entities from the article text x_i and types them with the retained SEMLEG classes. These class declarations are injected into the prompt as a compact Turtle Light serialization, following [8], to keep ontology guidance readable for the LLM. A second prompt π_i^{rel} then receives the article text, the extracted entities, and their classes, and generates triples T_i without a predefined predicate vocabulary. Relation labels are inferred from context, while subject and object classes restrict the plausible domain–range patterns. Finally, a third prompt π_i^{topic} assigns each triple to one of three topics: *maintenanceActivity*, *anotherLegalActivity*, or *legalCrossReference*. T_i is the resulting set of triples extracted from article a_i . The three prompt templates are available online.³

Embedding-Based Fusion of Entity and Property Labels. Triples classified as *maintenanceActivity* are retained for ontology induction T_{maint} ; triples assigned to the other topics are discarded at this stage. Since open extraction produces lexical variation, entity and property labels are normalized over this global set of maintenance triples. For each class of the ontology, we consider the entities which are its instances and occur as subject or object of the extracted triples. We embed their labels and merge them when their cosine similarity is greater than $\theta_E = 0.7$. The chosen canonical label is the most frequent variant. Legal references and numeric labels are excluded from this processing as they are unique or context-dependent values. Properties are normalized in the same way, grouping them according to the pairs of classes (c_{es}, c_{eo}) of which their subject and object are instances, and considering the same similarity threshold for merging their labels $\theta_P = 0.7$ [9, 10]. The canonical property label is selected as the most frequent variant. The resulting set T_{maint}^{fused} consists of the retained triples after replacing entity and property labels with their canonical forms.

Object Property Induction from Triples. To keep object property discovery tractable while preserving diversity, we sample in T_{maint}^{fused} 5% of triples per property and ensure at least one triple per signature. The sample is split into 15 batches. For each batch, a relation-discovery prompt $\pi_k^{\text{rel_disc}}$ receives the triples, the core ontology, and the competency questions, and generates object properties with label, domain, range, textual evidence, definition, aligned question, and confidence score. These property descriptions are formalized into OWL axioms.

Resulting Ontology. Two ontology variants are generated, one using OpenAI and one using Mistral for the generation of candidate object properties. Both reuse the same SEMLEG-based core classes and properties; as we can observe in Table 1, their differences come from the newly induced maintenance-specific object properties. The OpenAI variant produces a broader vocabulary, with 75 maintenance-

¹Légifrance platform: <https://www.legifrance.gouv.fr>.

²Apave regulatory guide: https://france.apave.com/Actualites/Publications/Livret_Obligations_reglementaires_Apave_2026.

³Code and prompts: <https://github.com/gmontenegrou/LegimaintLex>.

specific properties and 105 maintenance-specific signatures, while the Mistral variant is more compact, with 44 properties and 59 signatures. The two variants share 21 maintenance-related properties and 18 signatures, including `appliesTo`, `composedOf`, `performedAtLocation`, and `responsibleFor`, which capture recurring relations for applicability, composition, location, and responsibility.

Table 1

Comparison of the constructed SemLegM ontology variants.

	OpenAI	Mistral
# Maintenance-specific properties	75	44
# Maintenance-specific signatures	105	59

Qualitatively, OpenAI tends to induce more fine-grained and expressive predicates, including relations for purpose, sequencing, and interaction (e.g., `aimsToAction`, `precededBy`, `transmittedTo`). Mistral tends to produce a more conservative vocabulary oriented toward normative and compliance relations (e.g., `hasModality`, `verifiedBy`, `performedUnderCondition`). We keep the variants separate to compare how different ontologies affect downstream Knowledge Graph construction.

3. Automatic Construction of a French Legal Knowledge Graph

The constructed ontology variants are used for guiding closed large-scale triple extraction over the full corpus. Each ontology variant is paired with the LLM that generated it. In contrast to the ontology generation, where triple generation is open, the KG generation uses signature-level prompt guidance: the triple generation prompt receives the article text, the extracted entities, and ontology-derived constraints, including admissible object properties with their signatures. These constraints guide generation at prompt level, but are not enforced through a formal validation or repair loop.

After generation, the same fusion strategy is applied to reduce entity and predicate variation. The outputs are organized into intermediate tabular representations for legal documents, metadata, triples, and entity mentions, and then lifted to RDF using declarative RML mappings. The RDF model combines established vocabularies:⁴ ELI and DCTERMS for legal documents and metadata, CNT for textual content, SKOS and PROV for extracted entities and provenance, and Web Annotation for links between mentions and textual evidence. Relations are represented as reified statements typed as `rdf:Statement` and `semLeg:ExtractedRelation`, which makes it possible to attach evidence and contextual metadata to each extracted relation.

Table 2

Statistics of the knowledge graph variants with (F) or without (NF) fusion.

	Mistral NF	Mistral F	OpenAI NF	OpenAI F
# RDF triples	2,119,485	1,131,066	1,507,755	1,311,916
# Classes	18	15	15	12
# Entities	74,035	20,827	53,804	38,339
# Object properties	2,643	500	2,649	2,031
# Signatures	4,694	1,636	4,374	3,398
# <code>rdf:Statement</code> instances	75,870	75,870	51,657	51,657
# Annotations	119,063	39,055	83,879	70,930

Table 2 summarizes the size and variability of the generated graphs. The fusion clearly preserves the extracted relation statements while substantially reducing duplicated entities, predicates, signatures, and annotations, especially for the Mistral-based graph. This shows that normalization is a central step for turning raw LLM triple generation into a more compact and usable ontology-grounded KG.

⁴Vocabulary URLs: ELI <http://data.europa.eu/eli/ontology#>; DCTERMS <http://purl.org/dc/terms/>; CNT <http://www.w3.org/2011/content#>; SKOS <http://www.w3.org/2004/02/skos/core#>; PROV <http://www.w3.org/ns/prov#>; Web Annotation <http://www.w3.org/ns/oa#>.

4. Evaluation of the Ontology and Knowledge Graph

Experimental Setup Experiments were run with a Python pipeline using OpenAI (GPT-4.1 and text-embedding-3-large) and Mistral (mistral-large-2512 and mistral-embed-2312) under the same prompting protocol. Temperature was set to 0 for deterministic generation; we used a fixed context window size of 4k tokens for class-guided entity extraction and 10k for relation discovery. We evaluate both syntactic robustness and semantic alignment of the generated relations with the constructed ontology. Datasets are available online.⁵

Quantitative Analysis We use four compact metrics: JSON validity (R_{JSON}), class coverage (R_{class}), object property coverage (R_{prop}), and exact relation signature compliance (R_{sig}). R_{prop} measures whether a triple uses a property already present in the ontology, whereas R_{sig} additionally measures whether the property occurs with an expected domain–range combination. Table 3 shows robust structured output and near-complete class alignment. It also shows that fewer than 20% of triples introduce previously unseen properties in all settings, and much less after fusion. The lower R_{sig} values therefore point less to a lack of predicate vocabulary than to new or unexpected class combinations for existing predicates, such as frequent predicates reused across broader contexts than initially encoded in the ontology.

Table 3

Evaluation of legal relation statements with (F) or without (NF) fusion.

Metric	Mistral NF	Mistral F	OpenAI NF	OpenAI F
R_{JSON}	100.00%	100.00%	100.00%	100.00%
R_{class}	99.97%	99.97%	99.99%	99.99%
R_{sig}	49.86%	72.61%	56.94%	61.03%
R_{prop}	82.51%	96.42%	81.52%	85.11%

Qualitative Analysis We further inspect fused graphs by checking inconsistent relation signatures whose predicate suggests an expected target class (e.g., *hasTime* should point to *Time*). Mistral fusion yields 39 inconsistent signatures, while OpenAI fusion yields 52. Most true errors come from entity typing mistakes propagated to relations, such as legal sources classified as *Artifact*, or from predicates that encode an incorrect expected object type. In general, both models encode modality or polarity directly in predicates (e.g., *cannotApplyFor*, *mustNotExceed*), which increase variation and occasionally generate French predicates despite prompts explicitly requiring English labels. These cases show that predicate fusion must balance vocabulary unification with semantic precision: over-normalizing predicates may improve compactness, but can hide distinctions that are critical in industrial maintenance, where legal obligations, prohibitions, responsibilities, and conditions may imply operational risks. Competency-question tests, implemented as SPARQL queries, additionally confirm that the graph can retrieve actor roles and legal justifications for maintenance actions.

5. Conclusion and Future Work

We presented an LLM-assisted pipeline for constructing an ontology-grounded French legal knowledge graph for maintenance regulations. The pipeline extracts typed entities and open relations, normalizes labels through embedding-based fusion, induces candidate object properties from relation signatures, and uses the resulting ontology to guide KG construction. The evaluation results show strong structured-output validity and class alignment. They also show that most triples reuse properties already present in the ontology, with fewer than 20% involving previously unseen properties, while signature compliance reveals new domain–range combinations for existing predicates. Future work will therefore focus on iterative ontology refinement, not only by adding missing properties, but also by validating and

⁵Datasets: https://drive.google.com/drive/folders/1GfJPdJiGv_dZuC9utOgCTyzmeYb5lYU0?usp=drive_link.

integrating newly observed domain-range signatures. We will also improve entity and predicate fusion and explore formal validation mechanisms, e.g., using SHACL, to detect and repair invalid signatures. Finally, we will study how the generated KG can support GraphRAG-based consultation of maintenance regulations and how it can be updated when legal provisions evolve.

Acknowledgement. This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the project managed by the National Research Agency (ANR) with the reference number ANR-23-IACL-0001, and through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence Université Côte d’Azur under reference number ANR- 15-IDEX-01.

Declaration on Generative AI

During the preparation of this paper the authors used GPT-5.3 and Codex-5.5 for grammar and spelling check, paraphrase and reword, and formatting assistance. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication’s content.

References

- [1] D. Premasiri, T. Ranasinghe, R. Mitkov, M. El-Haj, I. Frommholz, Survey on legal information extraction: Current status and open challenges, *Knowledge and Information Systems* 67 (2025) 11287–11358. doi:10.1007/s10115-025-02600-5.
- [2] S. Li, L. Yi, A few-shot entity relation extraction method in the legal domain based on large language models, in: *Proceedings of the 2024 Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence*, ACM, Hong Kong, China, 2024, pp. 580–586. doi:10.1145/3675417.3675513.
- [3] J. Breton, M. B. Billami, M. Chevalier, C. Trojahn, Semantic model for the legal maintenance: The case of semantic annotation of france legislative and regulatory texts, in: *Workshop on Methodologies for Translating Legal Norms into Formal Representations (LN2FR 2022) in conjunction with JURIX 2022*, Saarbrücken, Germany, 2022. URL: <https://ut3-toulouseinp.hal.science/hal-04212544>.
- [4] J. Breton, M. B. Billami, M. Chevalier, C. Trojahn, Leveraging semantic model and LLM for bootstrapping a legal entity extraction: An industrial use case, in: *Knowledge Graphs in the Age of Language Models and Neuro-Symbolic AI: Proceedings of the 20th International Conference on Semantic Systems*, 17–19 September 2024, Amsterdam, The Netherlands, Studies on the Semantic Web, IOS Press, Amsterdam, Netherlands, 2024, pp. 20–36. doi:10.3233/ssw240004.
- [5] C. Woods, M. Selway, T. Bikaun, M. Stumtner, M. Hodkiewicz, An ontology for maintenance activities and its application to data quality, *Semantic Web* 15 (2024) 319–352. doi:10.3233/SW-233299.
- [6] V. A. Carriero, M. Scrocca, I. Baroni, A. Azzini, I. Celino, Procedural knowledge ontology (PKO), 2025. URL: <https://arxiv.org/abs/2503.20634v1>. doi:10.1007/978-3-031-94578-6_19.
- [7] M. Hodkiewicz, C. Woods, M. Selway, M. Stumtner, IOF-Maint: Modular maintenance ontology, 2024. URL: <http://arxiv.org/abs/2404.05224>. doi:10.26182/chzp-vs60.
- [8] C. Ringwald, F. Gandon, C. Faron, F. Michel, H. A. Akl, 12 shades of RDF: Impact of syntaxes on data extraction with language models, in: *Lecture Notes in Computer Science, Lecture Notes in Computer Science*, Springer Nature Switzerland, Hersanissos, Greece, 2024, pp. 81–91. doi:10.1007/978-3-031-78952-6_8.
- [9] Y. Lairgi, L. Moncla, R. Cazabet, K. Benabdeslem, P. Cléau, iText2KG: Incremental knowledge graphs construction using large language models, 2024. URL: <http://arxiv.org/abs/2409.03284>. doi:10.48550/arXiv.2409.03284.
- [10] Y. Lairgi, L. Moncla, K. Benabdeslem, R. Cazabet, P. Cléau, ATOM: Adaptive and optimized dynamic temporal knowledge graph construction using LLMs, 2025. URL: <http://arxiv.org/abs/2510.22590>. doi:10.48550/arXiv.2510.22590.