



A Unified Ontology-Based Knowledge Graph Framework for Multi-Domain Vietnamese Legal Document Retrieval

Vuong Thi Pham¹*, Minh Phan¹, Cao Bao Quynh¹, Truong Cong Phuc¹, Nguyen Thanh Duy¹, Tien Minh Vy¹, and Nguyen Ngoc Son¹

¹ Saigon University, Ho Chi Minh City, 700000, Vietnam.

KEYWORDS

Knowledge graph
Legal information retrieval
Multi-domain ontology
Natural language processing
Vietnamese legal documents

ABSTRACT

Legal information retrieval systems face the challenge of handling diverse legal domains with varying terminology, document structures, and semantic relationships. This paper presents a unified ontology-based knowledge graph framework that enables consistent knowledge representation and retrieval across multiple Vietnamese legal domains. The objective of this study is to develop a modular architecture that maximizes reuse while accommodating domain-specific requirements. Our methodology extends the Legal-Onto model with a three-layer architecture: Core Layer defining fundamental legal document structures based on Vietnamese legal formatting regulations, Domain Extension Layer providing specialized concept hierarchies for different legal areas, and Application Layer implementing query interfaces with natural language processing pipelines. The framework is validated through implementation and evaluation across three distinct legal domains: Land Law, Electronic Transaction Law, and Labor Law, comprising 53 legal documents with 1,398 concepts and 5,257 relationships. The results demonstrate that the unified framework achieves average accuracy of 84.2% and F1-score of 0.843 across all domains. Cross-domain evaluation reveals that domain-specific ontology extensions improve retrieval precision by 12.7% compared to generic approaches, while the shared core structure enables knowledge transfer and reduces development effort by 40% per new domain. In conclusion, the modular ontology architecture effectively balances reusability with domain specificity, providing a practical foundation for scalable legal information systems.

ARTICLE HISTORY

Received 31 December 2025
Received in revised form
28 May 2026
Accepted 11 June 2026
Available online 25 June 2026

© 2026 The Authors. Published by Penteract Technology.

This is an open access article under the CC BY-NC 4.0 license (<https://creativecommons.org/licenses/by-nc/4.0/>).

1. INTRODUCTION

The Vietnamese legal system comprises numerous domains, each with distinct terminology, document structures, and regulatory frameworks. From Land Law governing property rights to Electronic Transaction Law regulating digital commerce, and Labor Law protecting worker rights, each domain presents unique challenges for information retrieval systems. Traditional keyword-based search approaches struggle to capture the semantic relationships and domain-specific nuances essential for effective legal research [1], [2].

Knowledge graphs built on ontology foundations have emerged as powerful tools for representing structured domain knowledge [3], [4]. However, developing separate ontologies

and knowledge graphs for each legal domain is resource-intensive and fails to leverage commonalities across domains. Vietnamese legal documents share fundamental structural elements—chapters (Chương), sections (Mục), articles (Điều), clauses (Khoản), and points (Điểm)—that can form a unified representational framework following the structure defined in Circular 25/2011/TT-BTP [5], [6].

Recent advances in Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) have demonstrated promising results in legal question answering [7], [8]. However, these approaches still face challenges with hallucination and domain-specific accuracy that knowledge graph structures can help address [9], [10]. The combination of ontology-based knowledge representation with modern

*Corresponding author:

E-mail address: Vuong Thi Pham <vuong.pham@sgu.edu.vn>.

<https://doi.org/10.56532/mjsat.v6i2.684>

2785-8901/ © 2026 The Authors. Published by Penteract Technology.

This is an open access article under the CC BY-NC 4.0 license (<https://creativecommons.org/licenses/by-nc/4.0/>).

retrieval techniques offers a promising direction for robust legal information systems.

This paper addresses the challenge of building a unified ontology-based knowledge graph framework that supports multiple legal domains while accommodating domain-specific requirements. Our framework extends the Legal-Onto model [11] with a modular architecture comprising: (1) a shared core ontology capturing common legal document structures, (2) domain extension modules for specialized concepts and relationships, and (3) a unified query interface enabling cross-domain retrieval.

The main contributions of this paper include: (1) A modular ontology architecture that separates core legal structures from domain-specific extensions, enabling reuse and maintainability; (2) Comprehensive evaluation across three Vietnamese legal domains demonstrating the framework's versatility and effectiveness; (3) Quantitative analysis of cross-domain knowledge transfer and development efficiency gains of 40% per new domain.

Relation to previously published work. This manuscript is related to two previously published studies from the same broader research program. The first study presented an ontology-driven knowledge graph integrated with retrieval-augmented generation for Vietnamese legal question answering across Land Law, Electronic Transaction Law, and Labor Law [26]. The second study introduced a self-reflection mechanism for reducing hallucination in Vietnamese Labor Law question answering [27]. In contrast, the present paper focuses on the shared modular framework itself: it formalizes the three-layer Core/Domain/Application architecture, explains how domain extensions are attached to a reusable legal-ontology core, quantifies cross-domain reuse and development-efficiency gains, and reports a harmonized multi-domain evaluation at the framework level. Accordingly, the current manuscript should be read as the integrative architectural paper, whereas the prior publications describe application-oriented retrieval-generation and quality-control extensions built on top of related domain-specific components. Put differently, Reference [26] asks whether a knowledge-

graph-integrated RAG pipeline can deliver accurate, citation-grounded answers in Vietnamese legal question answering, Reference [27] asks whether self-reflection can reduce hallucination in the Labor Law setting, and the present paper asks whether a shared ontology architecture can support multiple legal domains efficiently.

2. METHODOLOGY

2.1. Related Work

Legal ontology development has a rich history, with foundational work by Sartor [12] establishing principles for structured legal knowledge representation. The LKIF ontology [13] provides a standardized vocabulary for legal concepts in European legal systems. For Vietnamese legal systems, Ngo et al. [5] developed an ontology knowledge map approach, while Pham et al. [11] created the Legal-Onto model for Vietnamese traffic law. Dang et al. [14] extended this work by integrating graph embeddings with ontology structures.

Multi-domain ontology integration remains an active research area. Hofer et al. [3] surveyed knowledge graph construction challenges, noting the tension between domain specificity and reusability. Modular ontology design patterns [15] offer strategies for managing complexity through layered architectures. For Vietnamese NLP infrastructure, PhoBERT [16] provides pre-trained language models, and VnCoreNLP [17] offers comprehensive text processing capabilities.

2.2. Framework Architecture

The unified Legal-Onto framework adopts a three-layer architecture designed to maximize reuse while accommodating domain variations (see Figure 1). The Core Layer defines fundamental ontology structure based on the Rela-model formalism, capturing universal Vietnamese legal document structures that apply across all domains. The Domain Extension Layer provides domain-specific concept hierarchies and relationships that extend the core with specialized legal knowledge. The Application Layer implements the query interface and NLP pipeline for end-user interaction.




Fig. 1. Unified Legal-Onto Framework Architecture showing three-layer design with Core Layer, Domain Extension Layer, and Application Layer

2.3. Core Ontology Structure

The core ontology follows the Legal-Onto formalism with enhanced modularity. Formally, the core knowledge base is defined as: $K_{core} = (C_{core}, R_{core}, Rules_{core})$, where C_{core} represents core concepts, R_{core} represents core relationships, and $Rules_{core}$ defines inference rules. Core concepts include structural elements common to all Vietnamese legal documents: Document (Văn bản), Chapter (Chương), Section (Mục), Article (Điều), Clause (Khoản), and Point (Điểm).

Core relationships capture the hierarchical structure: HAS_CHAPTER, HAS_SECTION, HAS_ARTICLE, HAS_CLAUSE, HAS_POINT for structural containment; REFERS_TO for cross-references; AMENDS and REPLACES for document versioning. These relationships apply uniformly across all legal domains.

To support legal version control, the AMENDS and REPLACES relations are modeled as temporal edges rather than ordinary semantic links. Each Document node stores version metadata including promulgation date, effective date, expiry date when available, and an is_current attribute. AMENDS connects an amending document or provision to the part of the earlier text it modifies, whereas REPLACES denotes full supersession by a later legal instrument. At query time, retrieval prioritizes nodes marked as current while preserving traversable links to historical versions for legal traceability and audit use.

2.4. Domain Extension Mechanism

Domain extensions inherit from core concepts while adding specialized attributes and relationships. Formally: $K_{domain} = K_{core} \cup (C_{ext}, R_{ext}, Rules_{ext})$, where the domain-specific extensions augment rather than override core definitions. For example, Land Law defines concepts like "Quyền sử dụng đất" (Land use rights); E-Transaction introduces "Chữ ký điện tử" (Electronic signature); Labor Law defines "Hợp đồng lao động" (Labor contract).

2.5. Multi-Domain Implementation

We implemented the framework across three Vietnamese legal domains: Land Law (Luật Đất đai 2024) [18] governing property rights, Electronic Transaction Law (Luật Giao dịch điện tử 2023) [19] regulating digital transactions, and Labor Law (Bộ luật Lao động 2019) [20] protecting worker rights. Each domain presents distinct characteristics in terminology and semantic complexity.

The present manuscript consolidates the reusable architectural elements that underlie our prior domain-specific implementations and downstream extensions [26], [27]. Its emphasis is therefore not on re-reporting a single application pipeline, but on identifying the stable core ontology, the domain-extension mechanism, and the framework-level reuse strategy that make multi-domain deployment feasible.




Fig. 2. Multi-Domain Knowledge Graph Comparison showing concept counts and relationship density across domains

2.6. Knowledge Graph Construction

For each domain, we followed a systematic construction process (see Figure 3): (1) Document parsing using regex patterns to extract structural elements; (2) Concept extraction using fine-tuned PhoBERT models [16] for Vietnamese legal named entity recognition; (3) Relationship extraction through pattern-based triplet generation; (4) Graph construction in Neo4j database with domain-specific node labels.

Before regex extraction, the parser performs text normalization for whitespace, numbering markers, punctuation, and line breaks to reduce noise introduced by PDF conversion and OCR. In practice, four categories of formatting deviation were observed beyond the canonical

Circular 25/2011/TT-BTP pattern: (1) pre-2011 documents using Roman numerals for section numbering; (2) implementing regulations that omit intermediate structural levels and proceed directly from Chapter to Article without Sections; (3) appended annexes formatted outside the main legal hierarchy; and (4) multi-document compilations that restart article numbering at each sub-document. The parser addresses these deviations through a normalization step followed by document-type-specific rule sets that select the most appropriate extraction patterns before hierarchical parsing.

In Neo4j, shared structural nodes retain common labels such as Document, Article, and Clause, while domain concepts are assigned both semantic labels and domain labels

together with canonical concept identifiers. This namespaced labeling scheme prevents cross-domain conceptual collision when similar lexical forms appear in multiple legal domains but refer to different legal meanings. At query time, domain isolation is enforced through Cypher query patterns that

constrain traversal to nodes carrying the target domain label. Cross-domain relationships such as REFERS_TO and AMENDS are resolved at the shared structural layer, allowing cross-document linkage without causing conceptual collision among domain-specific concepts.

Figure 4: Knowledge Graph Construction Pipeline




Fig. 3. Knowledge Graph Construction Pipeline: Document Parsing → Concept Extraction → Relationship Extraction → Graph Construction

3. RESULTS

3.1. Dataset Statistics

Table 1 presents the knowledge graph statistics across three legal domains. The total dataset comprises 53 legal documents with 1,032 articles, 1,398 concepts, 5,257 relationships, and 9,293 knowledge graph triplets.

Table 1. Knowledge graph statistics across three legal domains

Domain	Docs	Articles	Concepts	Relations	Triplets
Land Law	23	456	567	2,134	3,892
E-Transaction	12	187	342	1,247	2,156
Labor Law	18	389	489	1,876	3,245
Total	53	1,032	1,398	5,257	9,293

Note on corpus scope. To avoid ambiguity with the related application-oriented publication in Reference [26], the present framework validation uses a stable 53-document core corpus. Reference [26] reports a slightly broader 56-

document downstream question-answering configuration; the current manuscript retains the 53-document corpus in order to preserve a consistent architectural baseline across domains.

3.2. Retrieval Performance

Table 2 shows retrieval performance across domains. Test queries were collected from legal forums, university courses, and expert evaluations, totaling 180 queries per domain (540 total). The unified framework achieves consistent performance with average accuracy of 84.2% and F1-score of 0.843. All retrieval experiments in this section use the stable 53-document framework-validation corpus summarized in Table 1; Reference [26] reports a slightly broader downstream application configuration. The harmonized 540-query set used here was assembled for framework-level validation across the three domains; although some legal topics naturally overlap with downstream application settings in [26] and [27], the evaluation objective in the present paper is cross-domain architectural consistency rather than application-specific or hallucination-focused benchmarking.

Table 2. Retrieval performance across domains

Domain	Accuracy	Precision	Recall	F1-Score
Land Law	83.7%	85.2%	82.4%	0.838
E-Transaction	86.4%	88.1%	84.7%	0.864
Labor Law	82.5%	84.3%	80.9%	0.826
Average	84.2%	85.9%	82.7%	0.843




Fig. 4. Performance Metrics and Unified vs Standalone Comparison showing accuracy improvements across all domains

3.3. Ablation and Development Effort

To better explain the reported precision gain and implementation efficiency, we performed an ablation-style comparison between a generic core-only configuration and the full framework with domain-specific extensions. We also categorized implementation work into reusable core tasks and domain-extension tasks, enabling approximate person-hour comparison across domains.

The 40% reduction corresponds to $(200 - 120) / 200$ and was estimated by summing contributor-reported hours across the activity categories listed in Table 4.

Table 3. Ablation of generic core-only retrieval versus retrieval with domain-specific ontology extensions

Domain	Precision (Core-only)	Precision (Core + Domain Extension)	Gain	Main source of gain
Land Law	72.4%	85.2%	+12.8%	Land-right and parcel concepts reduced false matches
E-Transaction	75.4%	88.1%	+12.7%	Technical terminology and relation constraints narrowed candidate subgraphs
Labor Law	71.9%	84.3%	+12.4%	Contract and employment-status vocabulary improved query-to-concept alignment
Average	73.2%	85.9%	+12.7%	Domain lexicons and relation filters improved precision

Table 4. Person-hour estimation comparing standalone per-domain implementation and extension within the unified framework

Activity category	Standalone per-domain effort (hours)	Domain extension within unified framework (hours)
Ontology modeling and concept curation	52	28
Parser adaptation and normalization rules	28	12
Graph schema and Neo4j configuration	30	14
Query interface and retrieval adaptation	34	22
Domain lexicon, validation, and test queries	32	30
Integration, debugging, and refinement	24	14
Total	200	120

4. DISCUSSION

The unified framework consistently outperforms standalone implementations, with accuracy improvements of +1.6% (Land Law), +0.6% (E-Transaction), and +1.2% (Labor Law). These gains stem from the shared core ontology providing beneficial knowledge transfer across domains. Electronic Transaction Law achieves the highest performance (86.4% accuracy) because its concepts are expressed through relatively technical and standardized terminology, such as electronic signatures, certification services, and digital transaction procedures, which makes concept-relation matching more regular. Two additional factors likely contribute to this result. First, the E-Transaction corpus is the most compact among the three domains (12 documents, 187 articles), reducing the retrieval search space and the likelihood of false matches. Second, many questions in this domain were phrased in more formal and technically aligned language, which improves query-to-concept matching precision. Land Law is moderately more difficult because it contains denser cross-references and more context-dependent terms tied to land status, rights, and administrative procedures. Labor Law presents the most challenges (82.5% accuracy) because it covers a broader range of question types, including contracts, benefits, dismissal, obligations, and sanctions, and because users often formulate labor-law questions in colloquial or situation-based language rather than in canonical legal terms. Additionally, Labor Law provisions often involve multi-party conditional dependencies among employees, employers, trade unions, and state authorities, which require coordinated reasoning across multiple provisions. The current retrieval and ranking pipeline does not yet fully support such sustained dependency resolution, which further contributes to the lower accuracy observed in this domain.

Domain-specific ontology extensions improve retrieval precision by 12.7% compared to generic approaches without domain customization. As shown in Table 3, the gain is mainly attributable to three factors: domain-specific concept vocabularies reduce false matches between semantically adjacent but legally distinct terms, domain-specific relation constraints narrow the candidate subgraphs explored during retrieval, and specialized synonym mappings improve alignment between user queries and formal legal language. Development effort analysis reveals significant efficiency gains: the initial core framework required approximately 200 person-hours, while subsequent domain extensions required only 120 person-hours each on average, representing a 40% reduction compared to standalone implementations. These person-hours were estimated by summing contributor-reported effort across ontology design, parser adaptation, graph schema setup, query integration, validation, and refinement tasks, as summarized in Table 4.

Analysis of retrieval errors reveals three primary categories: (1) Ambiguous queries (38% of errors) where user intent spans multiple legal concepts; (2) Cross-reference gaps (32% of errors) involving provisions that reference decrees, circulars, or guidance documents not yet incorporated into the 53-document corpus; (3) Terminology variations (30% of errors) where informal language differs from formal legal terms. To reduce terminology-variation errors, the NLP pipeline can be extended with a domain synonym lexicon, a query normalization stage combining rule-based preprocessing with contextual paraphrase rewriting, and embedding-based

candidate expansion before graph matching. To address cross-reference gaps, a planned incremental expansion workflow consists of three steps: (1) constructing a reference dependency graph from REFERS_TO and AMENDS triples in the current corpus to identify frequently cited external documents; (2) retrieving those cited instruments from the official Vietnamese legal portal (vbpl.vn); and (3) triggering automated parsing and domain-extension updates for newly ingested documents. This strategy is intended to expand the corpus in proportion to actual reference demand rather than requiring full re-ingestion.

The present manuscript should also be distinguished from our recent application-oriented publications. Reference [26] emphasized citation-grounded integration of ontology-driven knowledge graphs with retrieval-augmented generation for Vietnamese legal question answering, while Reference [27] focused on a self-reflection layer for reducing hallucination in the Labor Law domain. Whereas [26] evaluates system-level answer accuracy and citation quality, and [27] evaluates hallucination reduction through iterative refinement, the present paper evaluates framework-level reusability and cross-domain development efficiency, which is a complementary but distinct research question. By contrast, this paper contributes the unified ontology architecture, the explicit separation between reusable core structures and domain-specific extensions, and the cross-domain efficiency analysis that explains how such downstream modules can be developed on a common foundation.

5. CONCLUSION

This paper presented a unified ontology-based knowledge graph framework for multi-domain Vietnamese legal document retrieval. By separating core legal structures from domain-specific extensions through a modular three-layer architecture, the framework achieves both reusability and domain specificity. The Core Layer captures universal Vietnamese legal document structures following Circular 25/2011/TT-BTP [6], while Domain Extension Layers provide specialized concept hierarchies for Land Law, Electronic Transaction Law, and Labor Law.

Experimental evaluation further demonstrates that the unified framework maintains consistent retrieval performance across all three domains, achieving an average accuracy of 84.2% and an F1-score of 0.843. Compared with standalone implementations, the unified approach improves accuracy by 0.6–1.6% and reduces development effort by 40% per new domain. In addition, domain-specific ontology extensions increase precision by 12.7% relative to a generic core-only configuration, confirming the practical value of combining reuse-oriented architecture with specialized domain modeling.

Future work will focus on four directions: (1) Extending the framework to additional Vietnamese legal domains including Criminal Law and Civil Law; (2) Implementing semi-automatic ontology extension mechanisms using LLM-based concept extraction [22], through a three-stage pipeline: candidate generation guided by the Core Layer schema, embedding-based ontology alignment to detect duplicates, and expert validation before integration; (3) Automatically expanding the corpus by detecting REFERS_TO links and ingesting cited decrees, circulars, and related legal instruments; and (4) Integrating the framework with RAG

systems for enhanced question answering capabilities [23]. This human-in-the-loop extension strategy is intended to reduce ontology drift while preserving the legal precision required in multi-domain retrieval systems.

In summary, the novel contribution of this manuscript is the framework-level unification itself rather than the isolated publication of downstream modules. This distinction is now made explicit to clarify the relationship between the current paper and the previously published extension studies [26], [27].

ACKNOWLEDGEMENT

This research is funded by Saigon University under Grant CSB.2025.87.

REFERENCES

- [1] A. Zalesinska et al., Legal Tech: Information Technology Tools in Justice Administration. Nomos Verlagsgesellschaft, 2021. doi: <https://doi.org/10.5771/9783748922834>
- [2] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," Advances in Neural Information Processing Systems, vol. 33, pp. 9459-9474, 2020. doi: <https://doi.org/10.48550/arXiv.2005.11401>
- [3] M. Hofer et al., "Construction of Knowledge Graphs: State and Challenges," Information, vol. 15, no. 8, pp. 1-25, 2024. doi: <https://doi.org/10.3390/info15080509>
- [4] A. Hogan et al., "Knowledge Graphs," ACM Computing Surveys, vol. 54, no. 4, pp. 1-37, 2022. doi: <https://doi.org/10.1145/3447772>
- [5] H. Q. Ngo et al., "Ontology Knowledge Map for Vietnamese Legal Applications," Vietnam Journal of Computer Science, vol. 11, no. 2, pp. 189-210, 2024. doi: <https://doi.org/10.1142/S2196888824500015>
- [6] Bộ Tư pháp Việt Nam, "Thông tư 25/2011/TT-BTP về thể thức và kỹ thuật trình bày văn bản quy phạm pháp luật," Hà Nội, 2011.
- [7] D. V. Dang et al., "Knowledge Graph-based Legal Query System with Large Language Models and RAG," Proc. Asian Conference on Intelligent Information and Database Systems, Springer, pp. 156-168, 2025. doi: https://doi.org/10.1007/978-981-97-4985-0_13
- [8] Y. Gao et al., "Retrieval-Augmented Generation for Large Language Models: A Survey," arXiv preprint arXiv:2312.10997, 2024. doi: <https://doi.org/10.48550/arXiv.2312.10997>
- [9] Z. Ji et al., "Survey of Hallucination in Natural Language Generation," ACM Computing Surveys, vol. 55, no. 12, pp. 1-38, 2023. doi: <https://doi.org/10.1145/3571730>
- [10] L. Huang et al., "A Survey on Hallucination in Large Language Models," arXiv preprint arXiv:2311.05232, 2023. doi: <https://doi.org/10.48550/arXiv.2311.05232>
- [11] V. T. Pham et al., "Ontology-based Intelligent Searching on Traffic Law Documents," Proc. International Conference on Agents and Artificial Intelligence, pp. 456-463, 2023. doi: <https://doi.org/10.5220/0011792600003393>
- [12] G. Sartor, "Approaches to Legal Ontologies," in Approaches to Legal Ontologies, Springer, pp. 1-14, 2011. doi: https://doi.org/10.1007/978-94-007-0120-5_1
- [13] R. Hoekstra et al., "The LKIF Core Ontology of Basic Legal Concepts," Proc. Workshop on Legal Ontologies and Artificial Intelligence Techniques, pp. 43-63, 2007.
- [14] D. Dang et al., "Information Retrieval Using Ontology and Graph Embeddings," Proc. Industrial, Engineering and Other Applications of Applied Intelligent Systems, Springer, pp. 287-299, 2023. doi: https://doi.org/10.1007/978-3-031-36822-6_25
- [15] A. Gangemi, "Ontology Design Patterns for Semantic Web Content," Proc. International Semantic Web Conference, Springer, pp. 262-276, 2005. doi: https://doi.org/10.1007/11574620_21
- [16] D. Q. Nguyen and A. T. Nguyen, "PhoBERT: Pre-trained Language Models for Vietnamese," Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 1037-1042, 2020. doi: <https://doi.org/10.18653/v1/2020.findings-emnlp.92>
- [17] T. S. Nguyen et al., "VnCoreNLP: A Vietnamese Natural Language Processing Toolkit," Proc. NAACL System Demonstrations, pp. 56-60, 2018. doi: <https://doi.org/10.18653/v1/N18-5012>
- [18] Quốc hội Việt Nam, "Luật Đất đai số 31/2024/QH15," Hà Nội, 2024.
- [19] Quốc hội Việt Nam, "Luật Giao dịch điện tử số 20/2023/QH15," Hà Nội, 2023.
- [20] Quốc hội Việt Nam, "Bộ luật Lao động số 45/2019/QH14," Hà Nội, 2019.
- [21] C. Yang et al., "Knowledge Graph and LLM Co-learning via Structure-oriented Retrieval-Augmented Generation," IEEE Data Engineering Bulletin, vol. 48, no. 1, pp. 45-58, 2025.
- [22] A. Madaan et al., "Self-Refine: Iterative Refinement with Self-Feedback," Advances in Neural Information Processing Systems, vol. 36, pp. 46534-46594, 2023. doi: <https://doi.org/10.48550/arXiv.2303.11761>
- [23] N. Shinn et al., "Reflexion: Language Agents with Verbal Reinforcement Learning," Advances in Neural Information Processing Systems, vol. 36, pp. 8634-8652, 2023. doi: <https://doi.org/10.48550/arXiv.2303.11366>
- [24] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," Proc. EMNLP-IJCNLP, pp. 3982-3992, 2019. doi: <https://doi.org/10.18653/v1/D19-1410>
- [25] S. Robertson and H. Zaragoza, "The Probabilistic Relevance Framework: BM25 and Beyond," Foundations and Trends in Information Retrieval, vol. 3, no. 4, pp. 333-389, 2020. doi: <https://doi.org/10.1561/15000000019>
- [26] V. T. Pham and M. Phan, "Integrating Knowledge Graph with Retrieval-Augmented Generation for Vietnamese Legal Question Answering," DS Journal of Digital Science and Technology, vol. 5, no. 1, pp. 1-13, 2026. doi: <https://doi.org/10.59232/DST-V511P101>
- [27] T. V. Pham et al., "A Self-Reflection Mechanism for Reducing Hallucination in Vietnamese Legal Question Answering Systems," Scientific Journal of Computer Science, vol. 2, no. 1, pp. 72-80, 2026. doi: <https://doi.org/10.64539/sjcs.v2i1.2026.380>