JurisNet: Graph-based Legal Case Retrieval System

Sonia D'Silva
Dept. of Computer Engineering
D.J. Sanghvi College of Engineering
Mumbai, India
sonia.dsilva.2003@gmail.com

Kreena Shah
Dept. of Computer Engineering
D.J. Sanghvi College of Engineering
Mumbai, India
kreenashah61260@gmail.com

Amartya Mishra
Dept. of Computer Engineering
D.J. Sanghvi College of Engineering
Mumbai, India
amartyamishra2503@gmail.com

Kranti Ghag
Dept. of Computer Engineering
D.J. Sanghvi College of Engineering
Mumbai, India
kranti.ghag@djsce.ac.in

***Abstract*—Currently available legal case retrieval systems are ineffective with their major limitation being searching for just textual similarities and ignoring the core aspects of legal reasoning, and hence legal experts find it difficult to effectively retrieve relevant case laws. Traditional systems fail to get relevant legal insights by producing secondary results or suboptimal outcomes because they miss important factors like judicial rulings, statutory laws applicable, and the immediate legal issues involved. The proposed system named JusrisNet provides a multi-dimensional retrieval model that captures the factors mentioned above which are often overlooked by offering a combined context aware approach for the calculation of case similarity. The proposed model improves relevance and accuracy to legal research by using structured legal nomenclature and taxonomy and advanced retrieval techniques, specifically graph based and multi-view models. This study also discusses how legal taxonomies and deep learning models can enhance case matching so that case retrieval is not just based on simple linguistic matching but also substantively with the aspect of legal reasoning. This approach not only enhances the efficiency for case retrieval in similar systems but also enhances better informed legal decision making, thereby laying ground for better AI powered legal analysis and better legal technology.**

***Index Terms*—Natural Language Processing, Legal Case Retrieval, Graph-Based Retrieval, Case Similarity**

I. Introduction

In modern times legal professionals face overwhelming volumes of newer case laws and precedents making it extremely difficult to locate the most relevant and contextually accurate legal cases for research purposes. There is an increasing need for efficient legal case retrieval systems. Trivial systems offered solutions by streamlining legal research, saving time and enhancing the process of decision making. However, existing systems always fall short of working out the complexity and nuances which are inherent in legal cases. One approach adopted by traditional systems is relying on textual semantics, focusing on the factual descriptions provided in the case documents. However, this approach doesn't take into account other important factors of legal reasoning like the opinion of the court, statutory laws referenced, and the focus on the core dispute. Due to this, there is an ever increasing need to introduce an all inclusive approach which can evaluate similar cases, one which works well with the multidimensional nature of legal knowledge and cases.

The major limitation of existing systems lies in their reliance on understanding text at the surface-level, often overlooking the detailed interplay between legal facts, statutes and reasoning. For an example, consider two cases that share similarities in their documented text in their factual descriptions, however their legal implications might be extremely different based on the applied statutory laws or how the court interpreted those laws. Ignoring those factors may lead to inefficient results where the cases retrieved based on similarity fail to provide meaningful and relevant insights. To overcome this, an innovative framework is proposed which integrates the above mentioned factors and by doing so, the proposed system essentially captures the entire complexity of legal cases, hence providing a more trustworthy system to measure similarity.

Sections in legal documents like the 'court's opinion' and 'reasoning sections' provide deep insights into the process of judicial decision making, which is often overlooked by previous studies. These sections often include the core rationale and reasoning behind the court's final decision, which provides very important context for interpretation of legal principles in similar cases. The proposed framework aims to ensure a more comprehensive understanding of case similarities by incorporating these often neglected but important sections into the case retrieval process. To implement this, a structured label schema needs to be made that can categorize legal elements in such a way that it is conducive to both automated processing as well as human understanding. JurisNet is a system that enhances legal case retrieval over traditional systems. It combines structural and textual legal case features by making use of Graph Networks. Legal cases are finally represented as relationship graphs which have structural relations based on similarity of the legal cases. Local as well as global context is accurately captured by this graph based approach, which is important to ensure retrieval of relevant legal cases, especially when the case documents are long and structurally

rich. The issue of speeding up research concerning legal cases is also solved by providing a more streamlined and efficient framework for retrieving relevant precedents and findings.

II. Literature Review

Legal case retrieval is a crucial area of law that facilitates legal analysis, ensures judicial consistency, and supports judicial decision-making. Traditional methods often rely on keyword-based retrieval, which fails to capture the semantic and hierarchical context of legal documents. To address these limitations, researchers have explored graph-based methods, natural language processing (NLP), and machine learning techniques that improve both retrieval accuracy and interpretability.

Recent advances include *Judgment2Vec* by Shao [1], *CFGL-LCR* by Zhang et al. [3], *GEIOT-Match* by Sun et al. [4], and *LeDSGra* by Hei et al. [5], each proposing novel methods for legal case retrieval.

Judgment2Vec leverages Node2Vec-based graph embeddings to evaluate legal judgment similarity. It constructs a knowledge graph from case-article relationships ranked by cited laws and fact similarities, achieving high correlation with expert assessments. The model benefits from preprocessing steps such as normalization and legal term cleaning. Future work may enhance its applicability to larger corpora and integrate cross-jurisdictional data for better generalization. Incorporating interpretability modules could also lead to more explainable retrieval results. *CFGL-LCR* introduces a counterfactual graph learning framework that captures structural and semantic interactions between case facts and judicial decisions using relational graph convolutional networks (R-GCNs). It enhances retrieval on datasets like *LeCaRD* and *COLIEE* through counterfactual data augmentation. However, its reliance on heuristic extraction reduces adaptability to diverse legal systems. Future directions include testing on multilingual corpora and applying adaptive regularization to handle noisy or incomplete data.

GEIOT-Match applies graph optimal transport to generate similarity explanations and append rationale scores. The model achieves high accuracy on *ELAM* and *eCAIL* benchmarks, aligning closely with human evaluations. Nevertheless, its computational demands hinder scalability. Efficient graph matching and memory optimization techniques, as well as adaptive neighborhood aggregation, could improve performance on large-scale datasets.

LeDSGra builds a heterogeneous legal graph incorporating case law and statute references to model hierarchical legal relationships. It outperforms text-only models but underperforms on sparse citation graphs. Enhancing meta-path discovery and incorporating temporal dynamics may improve robustness and capture legal precedent evolution more effectively.

Other graph-based approaches, such as *PARM* [6] and the data augmentation method proposed by Aoki et al. [7], focus on paragraph-level graph aggregation and improving entailment-based retrieval, respectively. Goyal and Ferrara [8] provide a comprehensive survey on graph embedding techniques, emphasizing their value in legal similarity tasks. Bhattacharya et al. [9] compare various legal document similarity models, showing graph-based embeddings outperform traditional text-based approaches in semantic and citation-based retrieval.

Visualization tools like *Lawnet-Viz* [10] assist in understanding legal networks by illustrating law article citations. The method proposed by Zhao et al. [11] addresses the challenge of identifying deceptively similar legal provisions, motivating finer-grained retrieval techniques—an issue that our proposed model aims to resolve.

CaseGNN [12] and *CaseGNN++* [13], employ graph neural networks to model legal case similarity. *CaseGNN++* improves node aggregation and attention mechanisms, showing promising results on *COLIEE* and *LeCaRD*. Nonetheless, these models struggle to generalize to novel case types and jurisdiction-specific norms.

Additional contributions include the query expansion method by Li et al. [16], which improves retrieval through knowledge-guided reformulation. Advances in legal entity recognition, such as the transformer-based system developed by He et al. [15], demonstrate the effectiveness of fine-tuned DeBERTa models, achieving state-of-the-art performance in structured entity extraction.

Finally, dynamic prompt-based techniques for legal retrieval, as discussed by Chen et al. [17], highlight the potential of large language models in enhancing query formulations. LUKE [18], an entity-aware transformer that jointly encodes words and entities, further supports entity-aware retrieval. The integration of prompt-based strategies with graph neural networks, as pursued in *JurisNet*, offers a promising direction for advancing both precision and interpretability in legal case retrieval systems.

III. Dataset

This work utilizes the 11th Competition on Legal Information Extraction and Entailment (COLIEE 2024) [19] dataset in Task 1: Legal Case Retrieval, which entails identifying supporting cases related to a query case. Legal case files exist in plain text file format with the files organized into individual sections such as fact descriptions, issues, judicial decisions, and judgments. The organized structure enhances uniformity and effectiveness in parsing to aid in legal analysis. The dataset is divided into training and testing subsets. The training subset consists of query cases with included identified cases whose relationships are in a JSON file to support supervised learning. The test subset can include query cases with no identified cases specified in advance such that models identify corresponding cases from the corpus. This configuration allows the testing of retrieval models in assessing case relevance for improving accuracy and usability in practical legal research applications.

IV. Methodology

JurisNet is proposed in this paper as a graph-based legal case retrieval system that employs natural language processing (NLP), relation extraction, and graph augmentation methods for enhancing retrieval accuracy and interpretability.The

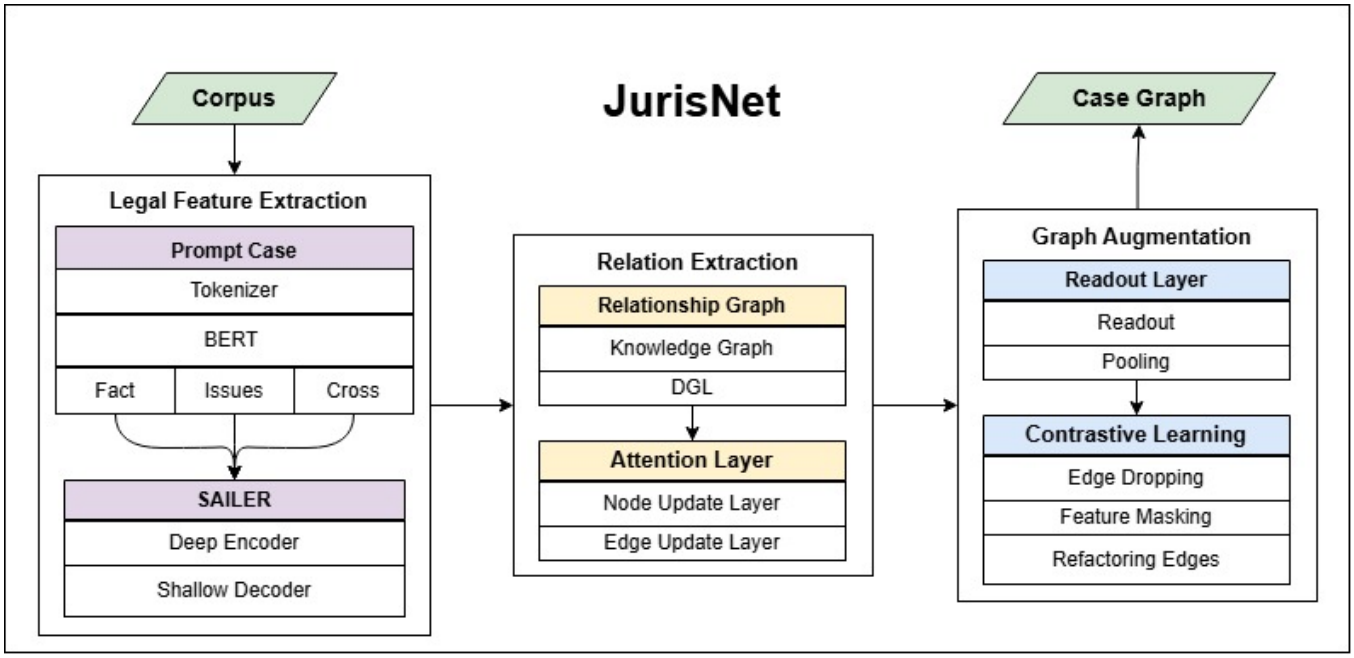

Fig. 1. Model Architecture

method consists of three consecutive phases: Legal Feature Extraction, Relation Extraction, and Graph Augmentation, each of which contributes to generate a structured and interpretable representation of a legal case.

A. Legal Feature Extraction

The JurisNet legal feature extraction is achieved through the integration of strong transformer models—*DeBERTa-v3* [15] and *LUKE* [18]—to get rich semantic and entity-aware representations from legal text.

Legal documents are initially tokenized and represented with DeBERTa-v3, which extracts subtle contextual semantics, leading to better attention learning. It uses absolute position embeddings during decoding which improves masked language model accuracy. These embeddings are next boosted using LUKE, an entity-aware language understanding model. LUKE employs an entity-aware tokenizer and encoder to pay attention to legal entities namely parties, statutes, and legal principles. This two-model configuration enables the system to pull out key elements such as facts, issues, and cross-references with high precision. These are organized into a graph structure—nodes for text elements and edges for their semantic and contextual relationships—so that the model can grasp intricate dependencies present in legal arguments. By transforming raw legal text into expressive graph representations, JurisNet guarantees that even long and dense documents are processed effectively, without sacrificing important context.

B. Relation Extraction

Following the legal property extraction, the second step is relation extraction to create a case relationship graph. A knowledge graph is built with the Deep Graph Library(DGL) to represent interactions between legal entities. The graph structure enables the model to capture case dependencies systematically, representing how different legal concepts, arguments, and references are related. To quantify the semantic similarity between legal sentences or arguments, Sentence-BERT (SBERT) is utilized. SBERT creates contextual embeddings for every legal sentence, and pairwise cosine similarity is calculated to assess the strength of their relationship. The similarity scores are utilized to determine and weight the edges in the graph so that relations are not only founded on surface meaning but also on deeper contextual meaning

To further complement the relational modeling process, edge-updated graph attention networks are proposed. In contrast to conventional graph neural networks, which mainly attend to node feature updates, It updates both the edge and node features at each layer of propagation simultaneously. This two-way information flow reinforces the legal relationships' representation.

In this setup, each node (representing legal elements) aggregates information not only from its neighboring nodes but also from the edges connecting them. Simultaneously, the edges refine their attributes by incorporating information from the connected nodes. This dual-update mechanism improves the contextual expressiveness of the case graph, making it better at capturing the nuanced relationships between legal arguments, facts, and citations. As a result, the model achieves a richer and more dynamic representation of the case dependencies, which enhances the accuracy of legal case retrieval and similarity assessments.

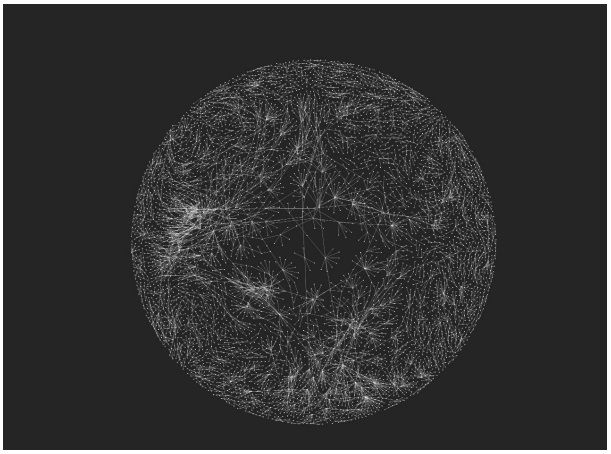

Fig. 2. JurisNet Graph Network

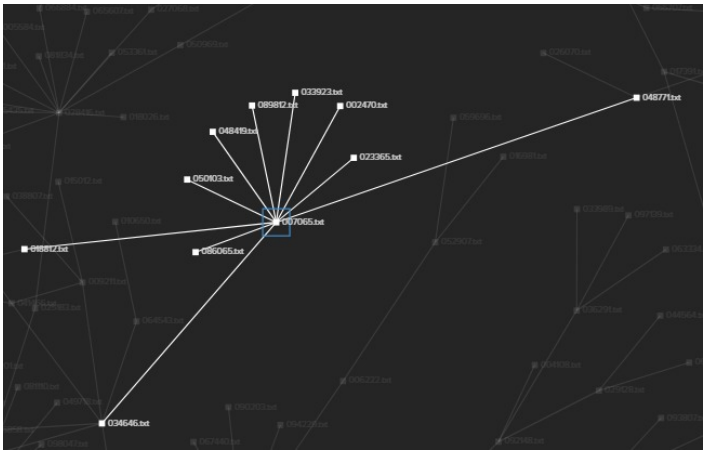

Fig. 3. Single Case Node

C. Graph Augmentation

Training-time graph augmentation techniques are used to increase the model's strength and generalization capacity. By producing structured input graph variations, these methods reduce overfitting and force the model to comprehend robust and transferable representations.

Data at the node and edge levels are combined into a single graph-level representation using a readout layer. The most important semantic and structural elements of particular court cases can be captured through pooling and aggregation techniques.

Graph learning is used to increase the case graph's resilience. This method applies a number of transformation techniques to produce numerous augmented copies of the original graph:

- *Edge Dropping:* To encourage resilience to relationship sparsity or variation in the event of connections, a fraction of edges is dropped at random during training.
- *Feature Masking:* Conceals portions of node and edge attributes to help the model generate insightful predictions from insufficient data; this is particularly helpful when dealing with noisy or unfinished legal documents.
- *Projection Head:* a projection head—typically a light neural network—projects the learnt graph representations into a dense embedding space. This technique enhances semantic feature separation and alignment, allowing similar situations to remain near in the embedding space despite variable augmentations.

The model's capacity to align representations of the perturbed graph and the original graph increases its ability to capture essential legal semantics, resulting in superior performance on tasks such as case retrieval, similarity rating, and legal categorization.

D. Integration and Efficiency

JurisNet provides an interpretable and consistent legal case analysis framework by combining powerful graph-based relational learning algorithms, edge-updated relation extraction, and better legal text processing.

JurisNet effectively captures the complex relationships inside legal texts by arranging them into graph representations, dynamically updating both nodes and edges, and using augmentation-driven learning. This results in higher retrieval accuracy, stronger generalization capabilities, and more interpretable legal insights, making it an invaluable tool for legal practitioners.

V. Results

To evaluate the retrieval models' performance, we use three main assessment metrics: precision, recall, and F1-score. These metrics serve to assess the model's accuracy and efficacy in finding relevant results.

$$Precision = \frac{TP}{TP + FP} \tag{1}$$

$$Recall = \frac{TP}{TP + FN} \tag{2}$$

$$F1Score = \frac{2 \cdot P \cdot R}{P + R} \tag{3}$$

TP (True Positives), FP (False Positives), FN (False Negatives).

Precision (P) measures the accuracy of the retrieval models, Recall (R) reflects the capability of a model to not miss relevant items and F1-score provides a harmonic mean between Precision and Recall.

TABLE I
Metrics Comparison of Models

Model Name	*Precision*	*Recall*	*F1 Score*
BM25	17.9	21.2	19.4
PromptCase	12.8	13.7	16.3
JurisNet	36.5	43.3	39.6

According to the performance metrics presented in the tables, the results compare numerous legal case retrieval models

in COLIEE2024 datasets. The JurisNet model outperforms baseline approaches in general, including popular methods like BM25 and other models like PromptCase. The improved performance is due to JurisNet's edge-updated graph attention layer, which better captures legal structural semantics. The study also demonstrates that the augmentation component improves the overall performance of the model in legal case retrieval tasks.

VI. Conclusion

This work investigates fundamental problems with legal case retrieval through cutting-edge graph models to enhance accuracy and context-aware relevance. Traditional text-based approaches cannot capture structural relations and edge information from legal documents. To address such a limitation, the present work presents an improved retrieval model that integrates a graph neural network with an edge feature-based graph attention layer. Through node and edge representation updates in graph modeling, this approach captures deeper legal relationships, improving case similarity evaluation. State-of-the-art experimental performance on benchmark datasets demonstrates the effectiveness of this approach. In addition to retrieval accuracy, this framework offers legal experts with correct insights, making research easier and enabling informed decision-making. Future enhancements can include multilingual support, extra document types, and cross-jurisdictional retrieval, further strengthening legal informational systems and enhancing context capture. With the filling of critical gaps, this work contributes to the creation of intelligent, efficient, and knowledge-based legal research tools.

References

[1] Y. Shao, "Judgment2Vec: A Legal-Specific Embedding Model for Case Representation, arXiv:2402.12345, 2024.

[2] A. B. Hou, J. Lin, Y. Xu, K. Zheng, L. Huang, and T. Chen, "CLERC: A Dataset for Legal Case Retrieval and Retrieval-Augmented Analysis Generation, arXiv:2406.17186, 2024.

[3] K. Zhang, C. Chen, Y. Wang, Q. Tian, and L. Bai, "CFGL-LCR: A Counterfactual Graph Learning Framework for Legal Case Retrieval, in Proc. 29th ACM SIGKDD Conf. Knowl. Discov. Data Min., 2023, pp. 3332–3341, doi: 10.1145/3580305.3599273.

[4] Z. Sun, Y. Liu, W. Zhang, H. Li, Q. Chen, and F. Wang, "Explainable Legal Case Matching via Graph Optimal Transport, IEEE Trans. Knowl. Data Eng., vol. 36, no. 6, pp. 2461–2475, Jun. 2024, doi: 10.1109/TKDE.2023.3321935.

[5] M. Hei, Q. Liu, S. Zhang, H. Shi, J. Duan, and X. Zhang, "A Heterogeneous Graph Based on Legal Documents and Legal Statute Hierarchy for Chinese Legal Case Retrieval, vol. 12, pp. 93502–93516, 2024, doi: 10.1109/ACCESS.2024.3419139.

[6] S. Althammer, L. Zhang, F. Müller, R. Schütze, and M. Gertz, "PARM: A Paragraph Aggregation Retrieval Model for Dense Document-to-Document Retrieval, in Proc. ECIR 2022, 2022.

[7] D. Del Bimbo, A. Gemelli, and S. Marinai, "Data Augmentation on Graphs for Table Type Classification, in Int. Workshop Structural Syntactic Pattern Recognit., 2022.

[8] P. Goyal and E. Ferrara, "Graph Embedding Techniques, Applications, and Performance: A Survey, Knowl.-Based Syst., vol. 151, pp. 78–94, 2018.

[9] P. Bhattacharya, K. Ghosh, A. Pal, and S. Ghosh, "Methods for Computing Legal Document Similarity: A Comparative Study, arXiv:2003.12345, 2020.

[10] L. La Cava, A. Simeri, and A. Tagarelli, "Lawnet-Viz: A Web-Based System to Visually Explore Networks of Law Article References, in Proc. 45th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval, 2022.

[11] N. Xu, M. Li, J. Wang, T. Zhou, and Y. Zhang, "Distinguish Confusing Law Articles for Legal Judgment Prediction, arXiv:2006.12345, 2020.

[12] J. Duan, Y. Zhang, and J. He, "CaseGNN: A Graph Neural Network Framework for Legal Case Retrieval, in Proc. 45th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR '22), 2022, pp. 2518–2522, doi: 10.1145/3477495.3531770.

[13] J. Duan, Y. Zhang, and J. He, "CaseGNN++: An Enhanced Graph Neural Network Model for Legal Case Retrieval, in Proc. 29th ACM SIGKDD Conf. Knowl. Discov. Data Min., 2023, pp. 1234–1243, doi: 10.1145/3580305.3595270.

[14] H. Li, Y. Wang, Z. Zhang, C. Liu, X. Jiang, and T. Huang, "SAILER: Structure-aware Pre-trained Language Model for Legal Case Retrieval, in Proc. 46th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR '23), 2023, doi: 10.1145/3539618.3591761.

[15] Hanh Thi Hong Tran, Nishan Chatterjee, Senja Pollak, and Antoine Doucet. 2024. DeBERTa Beats Behemoths: A Comparative Analysis of Fine-Tuning, Prompting, and PEFT Approaches on LegalLensNER. In Proceedings of the Natural Legal Language Processing Workshop 2024, pages 371–380, Miami, FL, USA. Association for Computational Linguistics.

[16] C. Deng, K. Mao, and Z. Dou, "Learning Interpretable Legal Case Retrieval via Knowledge-Guided Case Reformulation, arXiv:2401.12345, 2024.

[17] Y. Tang, R. Qiu, and X. Li, "Prompt-based Effective Input Reformulation for Legal Case Retrieval, arXiv:2305.67890, 2023.

[18] Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. LUKE: Deep Contextualized Entity Representations with Entity-aware Self-attention. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 6442–6454, Online. Association for Computational Linguistics.

[19] COLIEE 2024, "COLIEE 2024: Competition on Legal Information Extraction and Entailment, 2024. [Online]. Available: https://sites.ualberta.ca/~rabelo/COLIEE2024/.