Legal Knowledge and Information Systems
J. Savelka et al. (Eds.)

doi:10.3233/FAIA241283

LUMI: Legal Understanding and Matching Through Interactive Highlighting

Sabine WEHNERT [a,b,1], Visakh PADMANABHAN [b] and Ernesto William DE LUCA [a,b]

[a] *Leibniz Institute for Educational Media | Georg Eckert Institute, Brunswick, Germany*
[b] *Otto von Guericke University Magdeburg, Magdeburg, Germany*
ORCiD ID: Sabine Wehnert https://orcid.org/0000-0002-5290-0321, Ernesto William De Luca https://orcid.org/0000-0003-3621-4118

Abstract. In this paper, we present LUMI, a system that explains document retrieval through span highlighting. LUMI allows users to select a query span and highlights the most relevant part of a retrieved document using transformer-based retrieval, improving transparency in legal and technical analysis.

Keywords. Explainable Artificial Intelligence, Legal Information Retrieval, Legal Data Visualization

1. Introduction

Advancements in deep learning have improved Natural Language Processing (NLP), with transformers [1] and Large Language Models (LLMs), greatly impacting Information Retrieval (IR). While transformers perform well, their black-box nature limits explainability. This paper introduces LUMI[2], a novel tool using span highlighting to explain retrieval results, based on the Most Important Sentence [2], using transformer models such as SBERT [3]. LUMI focuses on identifying key substrings within documents that match user-selected query spans. To our knowledge, no other application offers this dynamic functionality, which is related to input perturbation techniques [4,5], where modifying parts of the input helps users understand which components drive the model's retrieval decisions. We demonstrate its use with the Task 3 dataset from Japanese civil law retrieval for the Competition on Legal Information Extraction/Entailment (COLIEE) [6], though it also has potential for various applications beyond legal data retrieval.

2. LUMI

LUMI helps legal experts identify which parts of a retrieved article best match a selected query segment, providing a similarity score for the match. The implementation is built

[1]Corresponding Author: Sabine Wehnert, sabine.wehnert@gei.de
[2]We share our code on github: `https://github.com/visakhpadmanabhan7/LUMI`

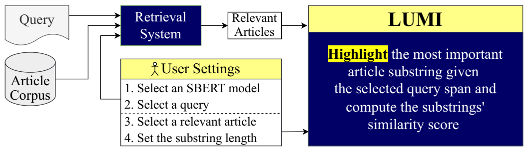

Figure 1. LUMI's Architecture

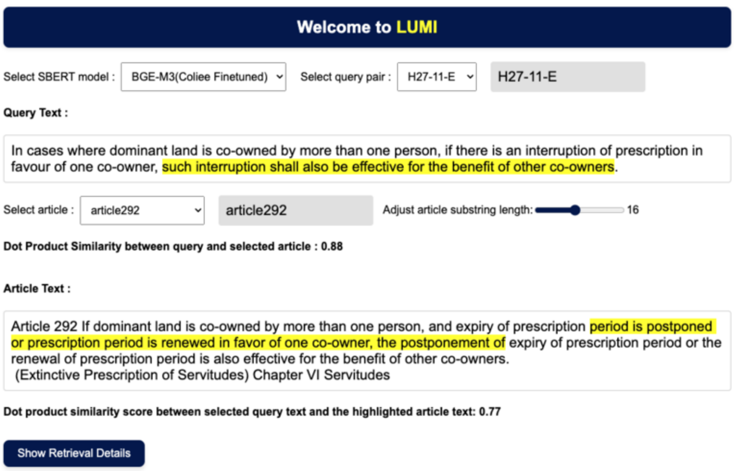

Figure 2. User Interface of LUMI

using Python Flask[3], while the user interface is developed with JavaScript. The overall process is illustrated in Figure 1. In the user interface (see Figure 2), the user first selects the SBERT model. During retrieval, we normalize the sentence embeddings from the SBERT model and calculate the dot product similarity between them (equal to cosine cosine similarity as embeddings are normalized). The dot product scores are then scaled between 0 and 1 using min-max scaling. A cutoff threshold, determined from previous experiments, is applied to provide candidate articles, in our use case from the whole set of all Japanese Civil Code articles that meet this threshold as relevant to the query, which we obtained from COLIEE Task 3 training data.

Once retrieval is complete, the user selects a query, chooses an article from the relevant candidates, and specifies a substring length to match within the article, given the highlighted span from the query (see the User Settings block in Figure 2). Finally, the user selects a portion of the query text, and the system highlights the most relevant substring from the article along with the corresponding similarity score. This approach ensures that the highlighted substring is both relevant and meets the similarity threshold. LUMI allows non-programmers to explore retrieval results and understand how the model identifies similarity and relevance in retrieved documents, by tracing how specific parts of documents align with query terms, e.g., to better understand false negatives.

[3] https://flask.palletsprojects.com/en/3.0.x/

References

[1] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is All you Need. In: Guyon I, von Luxburg U, Bengio S, Wallach HM, Fergus R, Vishwanathan SVN, et al., editors. Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA; 2017. p. 5998-6008. Available from: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

[2] Chen B, Chen K, Yang Y, Amini A, Saxena B, Chávez-García C, et al. Towards Improving the Explainability of Text-based Information Retrieval with Knowledge Graphs. CoRR. 2023;abs/2301.06974. Available from: https://doi.org/10.48550/arXiv.2301.06974.

[3] Reimers N, Gurevych I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In: Inui K, Jiang J, Ng V, Wan X, editors. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019. Association for Computational Linguistics; 2019. p. 3980-90. Available from: https://doi.org/10.18653/v1/D19-1410.

[4] Luo S, Ivison H, Han SC, Poon J. Local Interpretations for Explainable Natural Language Processing: A Survey. ACM Comput Surv. 2024;56(9):232:1-232:36. Available from: https://doi.org/10.1145/3649450.

[5] Sudhi V, Wehnert S, Homner NM, Ernst S, Gonter M, Krug A, et al. BiTe-REx: An Explainable Bilingual Text Retrieval System in the Automotive Domain. In: Amigó E, Castells P, Gonzalo J, Carterette B, Culpepper JS, Kazai G, editors. SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022. ACM; 2022. p. 3251-5. Available from: https://doi.org/10.1145/3477495.3531665.

[6] Goebel R, Kano Y, Kim MY, Rabelo J, Satoh K, Yoshioka M. Overview of Benchmark Datasets and Methods for the Legal Information Extraction/Entailment Competition (COLIEE) 2024. In: Suzumura T, Bono M, editors. New Frontiers in Artificial Intelligence. Singapore: Springer Nature Singapore; 2024. p. 109-24.