PATTERN-BASED PARSING OF GERMAN TRAFFIC REGULATIONS (STVO) FOR LEGAL KNOWLEDGE GRAPH CONSTRUCTION (KG_C)

Ibrahim Siddig, Sviatoslav Tugeev, Munir Georges

AImotion Bavaria, Technische Hochschule Ingolstadt, Germany
{Ibrahim.Siddig, Sviatoslav.Tugeev, Munir.Georges}@thi.de

Abstract: In this work, we present a domain-specific pattern-based parser (Π) designed to extract entities and relationships from the German traffic regulations legal text (**StVO**) using explicit citations and hierarchies within the text to construct a legal knowledge graph (*KG*). The domain knowledge embedded within the constructed knowledge graph (*KG*) is a valuable asset for a Large Language Model (LLM) training as part of future research in this domain. To the best of our knowledge, there has been no specific focus in the literature on parsing and tagging legal entities and their relationships for constructing knowledge graphs (*KGs*) from the German traffic regulations legal text, **Straßenverkehrs-Ordnung** (**StVO**), while focusing specifically on the hierarchical structure of the document. While other methods of extracting legal and factual knowledge from similar documents have been attempted, they either focused on different methodologies or different domains. The results indicate that the parser is capable of constructing a highly interconnected graph while maintaining the integrity of the original **Straßenverkehrs-Ordnung** (**StVO**) legal text, despite the graph's complexity.

1 Background

Utilization of traffic laws data contained within legal documents such as **Straßenverkehrs-Ordnung** (**StVO**) requires either formalization of the information contained within them or efficient processing of natural language. As was demonstrated by Sun et al. [1] and Esterle et al. [2], full formalization of traffic laws is difficult, both because they include highly context-dependent rules and because relevant information can be spread throughout multiple sections and requires efficient consolidation. Therefore, natural language processing remains as a potentially more viable approach for document interpretation. Existing LLMs have proven efficient in processing natural language documents [3]. Even non-specialized models are often capable of providing information related to traffic laws, without necessarily incorporating specific knowledge derived from legal documents of traffic regulations; unfortunately, the degree of imprecision that these models can demonstrate is incompatible with demands of the legal domain [4, 5].

Previous studies have demonstrated that retrieval-augmented information can significantly enhance performance during the inference phase [6]. Building on this, Srinivasagan and Georges [7] incorporate external knowledge to support inference generation, enabling the model to produce accurate information and generate high-quality, coherent sentences. Therefore, integrating domain-specific knowledge during the generation or training phase can improve the accuracy of such models [8, 6].

Knowledge graphs (*KGs*) have been shown to be a highly viable structure for enabling effecient information retrieval, among other tasks [9]. They offer a framework for connecting diverse data points, facilitating sophisticated data analysis [10]. Domain-agnostic knowledge

graph construction is possible [11], but it ignores the inherent relational and hierarchical structure already present within the document. As Winkels et al. [12] and Filtz [13] have demonstrated, relying on formal structure in the generation of knowledge graphs for legal documents specifically can be highly effective. Moreover, the accuracy of dependency parsing designed for a specific category of legal documents, such as traffic regulations, decreases when applied to different domains within this larger category, highlighting the need for specialized annotated corpora for effective domain adaptation [14]. Many knowledge-intensive applications (e.g. autonomous vehicles, drug discovery) rely heavily on the domain knowledge that is included in knowledge graph (*KGs*) ontologies [15].

Ontologies play a fundamental role in the organization and representation of information, providing structured frameworks for capturing the entities and their relationships within a domain and also serving as the basis for constructing knowledge graphs (*KGs*) [16]. Kommineni et al. [15] presents a semi-automated process that uses open-source LLMs to speed up the development of knowledge graphs (*KGs*) and ontologies. Although this method increases the accuracy of presented information, it is very dependent on prompt quality and faces difficulties like inconsistency, limited evaluation, and hallucinations [15]. Furthermore, domain-specific requirements, such as the demand for annotated datasets and expert knowledge, limit the scalability of LLM-based *KGs* production [15]. In the case of **StVO** it is therefore logical to focus on the top-down approach for *KG* generation [17], with predefined ontology based on the hierarchical and referential patterns within the document.

The difference between knowledge and data graphs is defined as the degree to which nodes and edges within the graph are organized with background knowledge about the concepts it represents [18, 19]. It is important to note that while the knowledge graph presented in this paper does not focus on extracting highly domain specific semantic relations and concepts from the **StVO** text, its ontology relies on structural hierarchy and categorization of **StVO** as a legal document, consisting of specific elements.

Section 1 established the background that motivated the choice of subject for this paper. Section 2 presents the employed methodology, outlining the structure of **StVO** document, proposed parser workflow and evaluation metrics used to measure its performance. Section 3 presents and analyses the results of this evaluation. Finally, Section 4 presents conclusions based on this analysis and outlines further potential avenues of research.

2 Methodology

This section will first briefly outline the structure of **StVO** text. It will then establish the proposed parser workflow, motivated by this structure. Finally, it will outline the chosen methods for evaluating the parser's performance.

2.1 Addressing Legal Documents of the German Traffic Rules

Legal documents in this domain contain semi-structured data, according to the classification proposed by Zhao et al. [20], and are typically divided into two distinct parts:

- **Main Law Text** (Part_1): includes a set of numbered **categories** (C) that contain **paragraphs** (P) that contain the essential legal content. Each paragraph discusses a particular subject and may contain additional subsections for explanations, such as **Absätze** (AB) and **Sätze** (S). In other words, a **paragraph** consists of a set of **Absätze**, and each **Absatz** is composed of a set of **Sätze**.

- **Supplementary Attachments** (Part_2): provide necessary legal references while adding more information to prevent overcrowding the core content. The main law **paragraphs** (P) may be expanded upon by diagrams, tables, or lists in this section, which are typically provided as appendices (A_n) listed as (Anlage 1, Anlage 2, etc.). Each appendix (A_n) includes components such as **Abschnitt** (CH), which further elaborate on additional elements like **Zeichen** (Z) as defined in the **StVO**.

2.2 StVO Parser Workflow

The parser (Π) is customized for the **StVO** using customized syntax patterns (Ψ) to parse each part (Part_1 and Part_2). The process follows a systematic, step-by-step approach, illustrated in Figure 1:

1. **Entities Extraction (E_X)**: Absätze (AB), categories (C), paragraphs (P), Sätze (sentences) (S), Abschnitt (CH), and Anlage (appendices) (A_n) represent the defined entities that the parser recognizes (E). These entities are extracted using hierarchical syntax patterns (Ψ) and contain the most relevant legal content.

2. **Building a Tree (T_B)**: Parent nodes (N_p) represent higher-level entities, child nodes (N_c) represent sub-elements, and leaves (N_l) indicate terminal entities in the hierarchical tree (T_B) into which the extracted entities are organized.

3. **Entity Tagging (T_G)**: Nodes (N) in the tree are tagged with their predefined entities (E), preserving explicit references and maintaining inter-entity connections.

4. **Entity Translation (T_R) (Optional)**: Since the original legal content is in German, the tagged entities (E) are translated into English using a translation pipeline (LLMs). This step ensures accessibility of the knowledge graph to non-German speakers and supports cross-lingual legal research.

5. **Knowledge Graph Construction (KG_C)**: The hierarchical structure (T_B) is transformed into a knowledge graph (KG) by defining edges (E_d) through explicit citations and syntactic relationships, such as parent-child connections (e.g., **paragraphs** (P) containing **Absätze** (AB)).

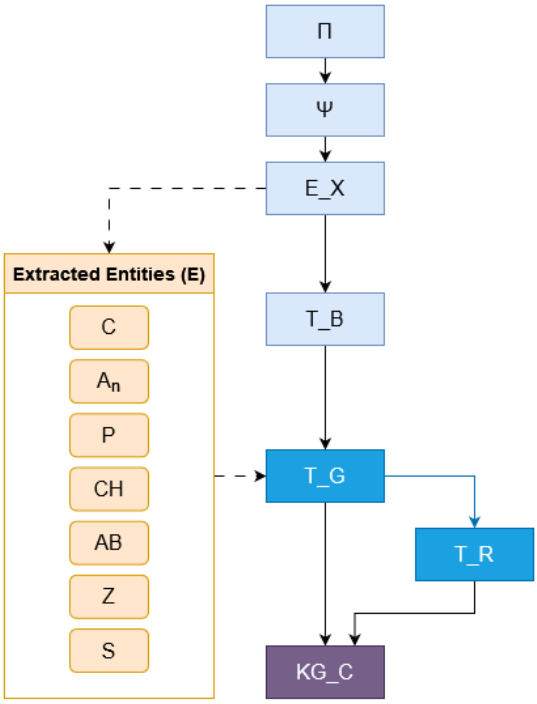

Figure 1 – The pipeline begins by parsing legal documents, specifically the German Road Traffic Regulations (StVO), in HTML format. Regular expression patterns are applied to extract entities defined in [2.2] such as paragraphs, sentences, and appendices, which are then organized hierarchically into a tree-based structure reflecting the relationships within the document. The nodes of this tree are tagged with their corresponding entities and optionally translated from German to English using large language models (LLMs) to enhance cross-lingual usability. Explicit citations within the tree nodes are leveraged to establish interconnections by adding edges from referencing nodes to referenced nodes. This step converts the hierarchical structure into a comprehensive knowledge graph (*KG*), supporting structured legal knowledge representation and cross-lingual accessibility. The final knowledge graph is shown in Figure 4.

2.3 Evaluation Metrics

The parser aims to extract and preserve the maximum amount of relevant legal information from the original text, while systematically structuring the it within a knowledge graph (*KG*). To evaluate the degree to which the original information is preserved after parsing, we attempt to use this information to accurately reconstruct the original legal text. Similarity in structure, content, and meaning between the original and reconstructed texts can therefore serve as a measure of the parser's performance. In order to fully measure the similarity between two texts, several key values need to be measured simultaneously. Consequently, several evaluation metrics were employed [21, 22, 23], each capturing a different aspect. The following metrics were used:

- **TF-IDF Similarity:** The textual similarity between the original and reconstructed texts is measured by the Term Frequency-Inverse Document Frequency (TF-IDF) similarity [21]. This metric calculates the cosine similarity between the two texts' TF-IDF vectors, reflecting how important a word is in a document relative to the entire corpus. It is particularly useful for evaluating the preservation of key terms and concepts in legal texts, as it emphasizes the importance of domain-specific vocabulary, indicating the preservation of the semantic essence of the original text by the parser. A similarity score between 0 (totally different) and 1 (identical) is the end result.

- **Word Error Rate (WER):** Quantifies the number of changes — substitutions, insertions, and deletions — required to transform the reconstructed text into the original text. [22] It is calculated as:

$$WER = \frac{\text{Insertions} + \text{Deletions} + \text{Substitutions}}{\text{Number of Words in Original (Reference) Text}} \tag{1}$$

  WER is a robust metric for assessing the parser's ability to reproduce the exact wording and structure of the original text, as it is highly sensitive to any modifications, including insertions, deletions, and substitutions. However, WER does not account for semantic similarity, meaning that rewording or synonym substitution is still counted as an error. This is particularly relevant for legal texts, where specific phrasing and sequence are of critical importance. A lower WER indicates a higher alignment between the texts, with a score of 0 representing a perfect match.

- **BLEU Score:** The Bilingual Evaluation Understudy (BLEU) score measures the n-gram overlap between the reconstructed text (candidate) and the original text (reference), making it useful for evaluating content preservation. The score ranges from 0 (no match) to 1 (perfect match) [23].

  The BLEU score complements WER by focusing on the overlap of phrases and sentences rather than individual words. This is important for assessing the parser's ability to maintain the coherence and flow of the original legal text.

These metrics were used together, as each of them works in tandem with the others to offer a comprehensive assessment of the parser's capacity to faithfully recreate the original legal text's content.

3 Results & Analysis

3.1 Performance of the StVO Parser

The parser's accuracy was evaluated by reconstructing the parsed content into plain text and comparing it against the original plain text using similarity metrics described in 2.3. Results of this comparison can be seen in Figure 2. TF-IDF Cosine Similarity score of 0.99, WER score of 0.12 and BLEU score of 0.86 reveal that the information extracted by the parser preserves the relevant aspects of the original text, specifically the phrasing, structure, and coherence.

Additional experiments in this study implement the parser on other traffic legal documents to test its generalizability. Generalizability, in this context, refers to the parser's ability to produce similar results when used on documents other than the one it was specifically designed for (i.e. the **StVO**). This includes handling variations in structure, formatting, and content while maintaining high representation accuracy. The experiments demonstrate that, while the parser is specifically designed for the **StVO**, it offers limited generalizability to other legal documents within the same domain (German traffic regulations). Limited generalizability here means that the parser can process and extract information from similarly structured documents but with reduced similarity scores between the original and the reconstructed texts, as indicated in Figure 3 compared to Figure 2. This reduction in accuracy could be attributed to the variability and unstructured nature of the **Supplementary Attachments** (Part_2) across traffic legal documents, which often include appendices, diagrams, or tables that differ significantly in format and content.

The parser described in this paper utilizes inherent hierarchical and reference structures of **StVO** to extract key legal information from the text and embed it into a knowledge graph. As **StVO** shares this highly formalized structure with many other categories of legal documents, the proposed parser could potentially be modified to effeciently extract key legal information from other similar documents. However, as can be seen in Figure 3, when applied to (**StVZO**), it produces higher WER and significantly lower BLEU scores, signifying that its ability to preserve the structure and coherence of the original text is currently highly dependant on the structural similarity of the processed documents to **StVO**.

3.2 Legal Knowledge Graph (*KG*)

Figure 4 illustrates a portion of the final knowledge graph (*KG*) constructed from the **StVO** using the parser workflow described in Section 2.2. Due to the large size of the full knowledge graph, the visualization in Figure 4 represents only a small subset of the entire structure. Specifically, the displayed subgraph is extracted by selecting a starting node and including only the nodes and edges that are reachable from this node.

In Table 1, we evaluate the structural complexity of the graph, which in this context refers to the extent of connectivity between nodes, including the number of nodes, edges, and the variation in branching factors. These metrics help assess how densely interconnected the graph is and how information propagates through its structure. A higher complexity often indicates richer relational data, allowing for more accurate query responses and improved knowledge retrieval. Conversely, extreme complexity can also introduce challenges in terms of computational efficiency and interpretability. The complexity is quantified using measurable metrics such as the average branching factor, which provides insights into the graph's organization and hierarchy. The large number of nodes (554) and outgoing edges (945) indicate a highly interconnected structure, contributing to the graph's overall complexity. The **average branching factor** of 1.71, calculated as:

$$\text{Average Branching Factor} = \frac{\text{Total Number of Outgoing Edges}}{\text{Total Number of Nodes}} \tag{2}$$

suggests that, on average, each node is connected to 1.71 edges. This value can be loosely related to the concept of **perplexity** in information theory, where a higher branching factor implies greater uncertainty or variability in the graph's structure. The gap between the minimum (0 edges) and maximum number of outgoing edges (46 edges) highlights the presence of both **terminal nodes** (e.g., child nodes like *Zeichen* nodes) and **hub nodes** (e.g., parent nodes like *paragraph* nodes). This disparity arises because terminal nodes represent endpoints in the hierarchy (e.g., individual Zeichen or sentences), while hub nodes act as central points connecting multiple sub-elements (e.g., paragraphs containing multiple *Absätze*).

Metric	**Value**
Number of Nodes	554
Number of Outgoing Edges	945
Average Branching Factor	**1.706**
Minimum Out-Degree	0
Maximum Out-Degree	46

Table 1 – Evaluation metrics of the graph showcasing its structural interconnectivity, including the number of nodes, outgoing edges, and degree variations. The **Average Branching Factor** is highlighted as a key measure of node connectivity.

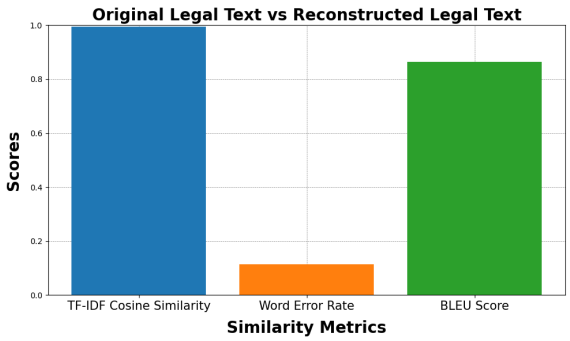

Figure 2 – Evaluation of parser accuracy through document similarity by comparing the reconstructed content from the parsed document with the original content Straßenverkehrs-Ordnung **(StVO)** using similarity metrics [2.3].
TF-IDF Similarity: 0.99; **WER**: 0.12; **BLEU**: 0.86.

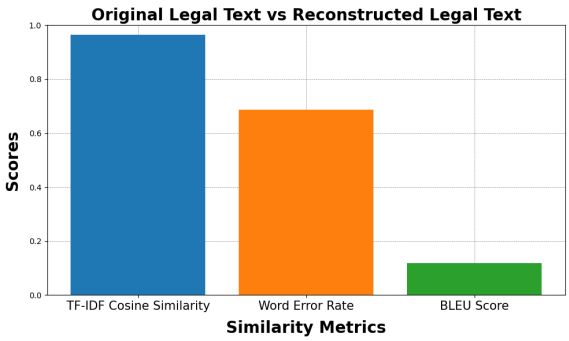

Figure 3 – Evaluation of parser accuracy through document similarity by comparing the reconstructed content from the parsed document with the original content Straßenverkehrs-Zulassungs-Ordnung **(StVZO)** using similarity metrics [2.3].
TF-IDF Similarity: 0.96; **WER**: 0.69; **BLEU**: 0.12.

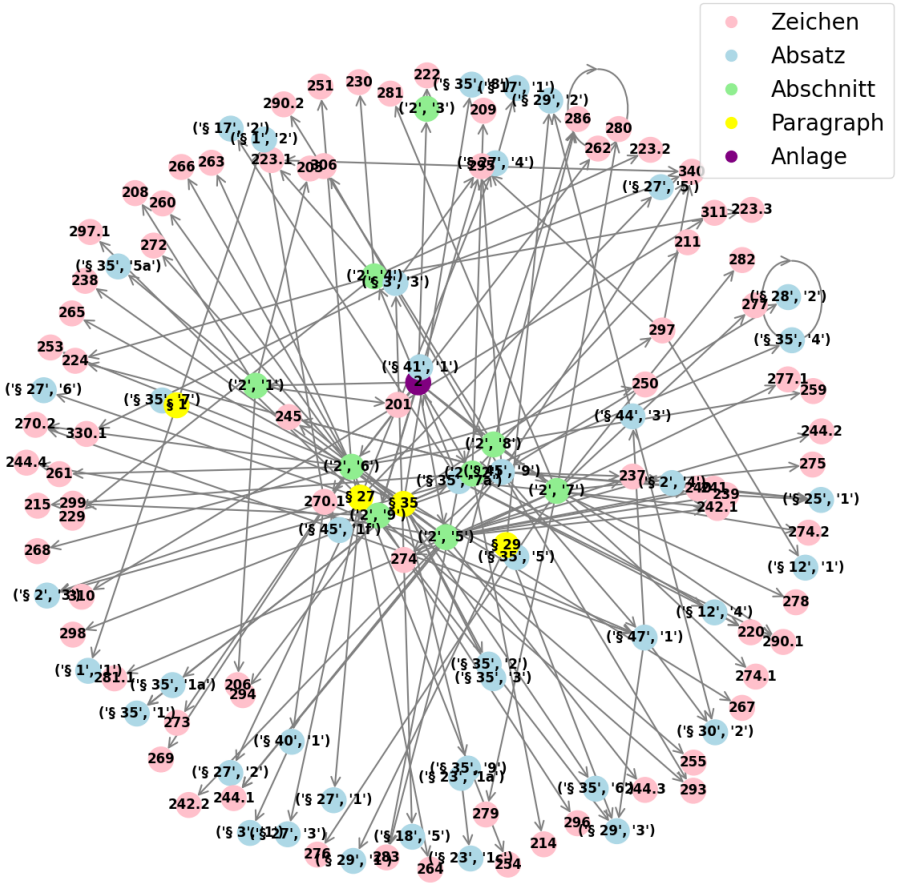

Figure 4 – A portion of the knowledge graph (*KG*) that was created using the suggested parsing framework in Figure 1 and derived from the Straßenverkehrs-Ordnung (StVO). Every node has tuples that correspond to its unique IDs, and the node content incorporates the related legal knowledge.

4 Conclusion

The approach described in this paper focuses on constructing knowledge graphs (*KGs*) based on the formal hierarchical and referential structure of the document. When applied to **StVO**, it demonstrates high RF-IDF and BLEU scores and low WER score, indicating that the coherence, meaning and structure of the original legal text are preserved in the constructed knowledge graph (*KG*). This knowledge graph could be highly beneficial for increasing the accuracy of LLMs when working with traffic laws and **StVO** specifically.

Proposed parser workflow is particularly well suited for highly structured legal documents with explicit and hierarchical relationships between elements, but, as seen in Figure 3, there are limits to its generalizability. Exploring and expanding these limits is a promising future avenue of research. Additionally, to capture more complex non-structural semantic relationships within the document and to expand the graph with additional domain-specific entities, integration of annotated datasets along with NLP-based dependency parsers, such as one described by Sun et al. [24] should be further investigated as a possible avenue of future research as well.

References

[1] SUN, Y., C. M. POSKITT, J. SUN, Y. CHEN, and Z. YANG: *LawBreaker: An Approach for Specifying Traffic Laws and Fuzzing Autonomous Vehicles*. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, ASE '22, pp. 1–12. Association for Computing Machinery, New York, NY, USA, 2023. doi:10.1145/3551349.3556897. URL https://dl.acm.org/doi/10.1145/3551349.3556897.

[2] ESTERLE, K., L. GRESSENBUCH, and A. KNOLL: *Formalizing Traffic Rules for Machine Interpretability*. In *2020 IEEE 3rd Connected and Automated Vehicles Symposium (CAVS)*, pp. 1–7. 2020. doi:10.1109/CAVS51000.2020.9334599. URL https://ieeexplore.ieee.org/abstract/document/9334599.

[3] TSIRMPAS, D., I. GKIONIS, G. T. PAPADOPOULOS, and I. MADEMLIS: *Neural natural language processing for long texts: A survey on classification and summarization*. *Engineering Applications of Artificial Intelligence*, 133, p. 108231, 2024. doi:10.1016/j.engappai.2024.108231. URL https://www.sciencedirect.com/science/article/pii/S0952197624003890.

[4] CHEONG, I., K. XIA, K. J. K. FENG, Q. Z. CHEN, and A. X. ZHANG: *(A)I Am Not a Lawyer, But...: Engaging Legal Experts towards Responsible LLM Policies for Legal Advice*. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, pp. 2454–2469. Association for Computing Machinery, New York, NY, USA, 2024. doi:10.1145/3630106.3659048. URL https://dl.acm.org/doi/10.1145/3630106.3659048.

[5] SHUI, R., Y. CAO, X. WANG, and T.-S. CHUA: *A Comprehensive Evaluation of Large Language Models on Legal Judgment Prediction*. In H. BOUAMOR, J. PINO, and K. BALI (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 7337–7348. Association for Computational Linguistics, Singapore, 2023. doi:10.18653/v1/2023.findings-emnlp.490. URL https://aclanthology.org/2023.findings-emnlp.490/.

[6] WANG, X., Q. YANG, Y. QIU, J. LIANG, Q. HE, Z. GU, Y. XIAO, and W. WANG: *Knowledgpt: Enhancing large language models with retrieval and storage access on knowledge bases*. *arXiv preprint arXiv:2308.11761*, 2023.

[7] SRINIVASAGAN, G. and M. GEORGES: *Retrieval augmented spoken language generation for transport domain*. In E. NÖTH, A. HORÁK, and P. SOJKA (eds.), *Text, Speech, and Dialogue: 27th International Conference, TSD 2024 Brno, Czech Republic, September 9-13, 2024 Proceedings, Part II*, pp. 3 – 12. Springer, Cham, 2024. doi:https://doi.org/10.1007/978-3-031-70566-3_1.

[8] ASAI, A., Z. WU, Y. WANG, A. SIL, and H. HAJISHIRZI: *Self-rag: Learning to retrieve, generate, and critique through self-reflection. arXiv preprint arXiv:2310.11511*, 2023.

[9] WU, D. and M. GEORGES: *A pipeline for automatic construction and applications of college curriculum knowledge graph*. In *TSD 2024 Paper Abstracts*. 2024. URL `https://www.tsdconference.org/tsd2024/abstracts.html#I1270`.

[10] PENG, C., F. XIA, M. NASERIPARSA, and F. OSBORNE: *Knowledge Graphs: Opportunities and Challenges*. *Artificial Intelligence Review*, 56(11), pp. 13071–13102, 2023. doi:10.1007/s10462-023-10465-9. URL `https://doi.org/10.1007/s10462-023-10465-9`.

[11] HAN, Y., W. DENG, Z. WANG, and K. WANG: *An Improved Method for Constructing Domain-Agnostic Knowledge Graphs*. In A. EL ABBADI, G. DOBBIE, Z. FENG, L. CHEN, X. TAO, Y. SHAO, and H. YIN (eds.), *Database Systems for Advanced Applications. DASFAA 2023 International Workshops*, pp. 222–237. Springer Nature Switzerland, Cham, 2023. doi:10.1007/978-3-031-35415-1_16.

[12] WINKELS, R., A. BOER, E. DE MAAT, T. VAN ENGERS, M. BREEBAART, and H. MELGER: *Constructing a semantic network for legal content*. In *Proceedings of the 10th international conference on Artificial intelligence and law*, ICAIL '05, pp. 125–132. Association for Computing Machinery, New York, NY, USA, 2005. doi:10.1145/1165485.1165505. URL `https://dl.acm.org/doi/10.1145/1165485.1165505`.

[13] FILTZ, E.: *Building and Processing a Knowledge-Graph for Legal Data*. In E. BLOMQVIST, D. MAYNARD, A. GANGEMI, R. HOEKSTRA, P. HITZLER, and O. HARTIG (eds.), *The Semantic Web*, pp. 184–194. Springer International Publishing, Cham, 2017. doi:10.1007/978-3-319-58451-5_13.

[14] DELL'ORLETTA, F., S. MARCHI, S. MONTEMAGNI, B. PLANK, and G. VENTURI: *The splet-2012 shared task on dependency parsing of legal texts*. In *Proceedings of the 4th Workshop on Semantic Processing of Legal Texts*, vol. 4. 2012.

[15] KOMMINENI, V. K., B. KÖNIG-RIES, and S. SAMUEL: *From human experts to machines: An llm supported approach to ontology and knowledge graph construction. arXiv preprint arXiv:2403.08345*, 2024.

[16] HOGAN, A., E. BLOMQVIST, M. COCHEZ, C. D'AMATO, G. D. MELO, C. GUTIERREZ, S. KIRRANE, J. E. L. GAYO, R. NAVIGLI, S. NEUMAIER ET AL.: *Knowledge graphs*. *ACM Computing Surveys (Csur)*, 54(4), pp. 1–37, 2021.

[17] LI, Y., H. YIN, R. LI, and H. LI: *A Knowledge Graph Construction Method for Legal Documents*. In Y. JIA, W. ZHANG, Y. FU, and H. YANG (eds.), *Proceedings of 2024 Chinese Intelligent Systems Conference*, vol. 1283, pp. 62–70. Springer Nature Singapore, Singapore, 2024. doi:10.1007/978-981-97-8650-3_7. URL https://link.springer.com/10.1007/978-981-97-8650-3_7.

[18] ZHONG, L., J. WU, Q. LI, H. PENG, and X. WU: *A Comprehensive Survey on Automatic Knowledge Graph Construction*. *Acm Computing Surveys*, 56(4), pp. 94:1–94:62, 2023. doi:10.1145/3618295. URL https://dl.acm.org/doi/10.1145/3618295.

[19] WU, M. and X. WU: *On big wisdom*. *Knowledge and Information Systems*, 58(1), pp. 1–8, 2019. doi:10.1007/s10115-018-1282-y. URL https://doi.org/10.1007/s10115-018-1282-y.

[20] PONTIKI, M., D. GALANIS, H. PAPAGEORGIOU, I. ANDROUTSOPOULOS, S. MANANDHAR, M. AL-SMADI, M. AL-AYYOUB, Y. ZHAO, B. QIN, O. DE CLERCQ ET AL.: *Semeval-2016 task 5: aspect based sentiment analysis (. pdf)*. *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, 2016.

[21] AIZAWA, A.: *An information-theoretic perspective of tf–idf measures*. *Information Processing & Management*, 39(1), pp. 45–65, 2003.

[22] ALI, A. and S. RENALS: *Word error rate estimation for speech recognition: e-WER*. In I. GUREVYCH and Y. MIYAO (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 20–24. Association for Computational Linguistics, Melbourne, Australia, 2018. doi:10.18653/v1/P18-2004. URL https://aclanthology.org/P18-2004/.

[23] PAPINENI, K., S. ROUKOS, T. WARD, and W.-J. ZHU: *Bleu: a method for automatic evaluation of machine translation*. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318. 2002.

[24] SUN, Q., Y. LUO, W. ZHANG, S. LI, J. LI, K. NIU, X. KONG, and W. LIU: *Docs2KG: Unified Knowledge Graph Construction from Heterogeneous Documents Assisted by Large Language Models*. 2024. doi:10.48550/arXiv.2406.02962. URL http://arxiv.org/abs/2406.02962.