

uOttawa at LegalLens-2024: Transformer-based Classification Experiments

Nima Meghdadi and Diana Inkpen

School of Electrical Engineering and Computer Science
University of Ottawa,
Ottawa, ON, K1N 6N5

nmegh082@uottawa.ca diana.inkpen@uottawa.ca

Abstract

This paper presents the methods used for LegalLens-2024 shared task, which focused on detecting legal violations within unstructured textual data and associating these violations with potentially affected individuals. The shared task included two subtasks: A) Legal Named Entity Recognition (L-NER) and B) Legal Natural Language Inference (L-NLI). For subtask A, we utilized the spaCy library, while for subtask B, we employed a combined model incorporating RoBERTa and CNN. Our results were 86.3% in the L-NER subtask and 88.25% in the L-NLI subtask. Overall, our paper demonstrates the effectiveness of transformer models in addressing complex tasks in the legal domain. The source code for our implementation is publicly available at <https://github.com/NimaMeghdadi/uOttawa-at-LegalLens-2024-Transformer-based-Classification>

1 Introduction

The huge amount of information and massive use of the internet has propelled to ignore legal violations, individual rights, cultural values and societal norms. These hidden violations demand serious attention and urgent solution due to serious effects on rights and justice and it requires advanced tools for professionals to effectively manage large amount of paperwork.

Legal violation identification seeks to automatically detect legal violations within unstructured text and link these violations to potential victims. The LegalLens 2024 shared task (Bernsohn et al., 2024) aims to foster a legal research community by tackling two key challenges in the legal domain. Subtask A focuses on identifying legal violations (a.k.a Identification Setup) using Named Entity Recognition (NER). Subtask B focuses on linking these violations to potentially affected individuals (a.k.a Identification Setup) using Natural Language Inference (NLI).

Our team participated in both subtasks of the shared task. In subtask A, we used the spaCy library and a DeBERTa-based model. In subtask B, we developed a RoBERTa-based model combined with a CNN-based model.

2 Related Work

There has been extensive research on Legal Named Entity Recognition (NER) for German legal documents. Leitner et al. (2019) developed NER models using CRF and BiLSTM, while Darji et al. (2023) used a BERT-based model. Many languages are using NER to expedite the process of judicial decision-making. For the Turkish language, Çetindağ et al. (2023) developed an NER model using BiLSTM and several word embeddings like GloVe, Morph2Vec, and neural network-based character feature extraction techniques. In Portuguese, Bonifacio et al. (2020) and Albuquerque et al. (2023) focused on NER models specific to the legal domain. The former developed a model using ELMo and BERT with the LeNER-Br dataset (Luz de Araujo et al., 2018), while the latter evaluated BiLSTM+CRF and fine-tuned BERT models on legal and legislative domains to automate and accelerate tasks such as analysis, categorization, search, and summarization. In Italian, Pozzi et al. (2023) created a model that combines transformer-based Named Entity Recognition (NER), transformer-based Named Entity Linking (NEL), and NIL prediction. In Chinese, Zhang et al. (2023) proposed a NER method for the legal domain named RoBERTa-GlobalPointer, combining character-level and word-level feature representations using RoBERTa and Skip-Gram, which were then concatenated and scored with the GlobalPointer method. Lee et al. (2023) also developed a legal domain NER model called LeArNER, which employs Bouma’s unsupervised learning for feature extraction and utilizes the LERT and LSTM

models for sequence annotation.

Kim et al. (2024) described methods for the COLIEE 2023 competition, using a sentence transformer model for case law retrieval and a fine-tuned DeBERTa model for legal entailment that used SNLI (Bowman et al., 2015) and MultiNLI (Williams et al., 2018) datasets for training. Tang (2023) explored improving legal Natural Language Inference (NLI) by employing general NLI datasets with supervised fine-tuning and examining the impact of transfer learning from Adversarial NLI to ContractNLI. The objective of Valentino and Freitas (2024) is to offer a theoretically grounded characterization of explanation-based Natural Language Inference (NLI) by integrating contemporary philosophical accounts of scientific explanation with an analysis of natural language explanation corpora. Gubelmann et al. (2023) investigated how large language models (LLMs) handle different pragmatic sentence types, like questions and commands, in natural language inference (NLI), highlighting the insensitivity of MNLI and its fine-tuned models to these sentence types. It developed and publicly released fine-tuning datasets to address this issue and explored ChatGPT’s approach to entailment.

3 Subtask A: Legal Named Entity Recognition(L-NER)

Subtask A, which involves finding named entities for specific types that may appear in legal texts, is explained in this section.

We developed a BERT-based model for this subtask as part of the LegalLens task, achieving an F1-score(Macro F1-score) of 86.3%.

3.1 Dataset Details

We used the dataset provided by the organizers of the shared task. The provided data was split into training and test sets, with each set consisting of tokenized text and the corresponding entities for those tokens. It is important to note that the provided test set includes labeled data, which is different from the separate test data that the organizers will use to evaluate the model. The split dataset used for validation in this research consists of 20% of training data and is shown in Table 1.

The entity types are fully described in (Bernsohn et al., 2024). The labels include four entity types: violation, violation by, violation on, and law, with detailed counts for each entity available in (Bern-

Type	Number of documents
Training	568
Validation	142
Test	617

Table 1: The number of documents used to train the model is detailed

Hyperparameter	Value
Learning Rate	5e-5
Batch Size	16
Maximum Steps	20,000
Dropout Rate	0.1
Optimizer	Adam

Table 2: Hyperparameters of the fine-tuned model for subtask A (L-NER)

sohn et al., 2024).

3.2 Preprocessing

For this subtask, we configure the spaCy pipeline with an emphasis on tokenization and vector initialization. The tokenizer used is the standard spaCy tokenizer, which splits the text into tokens for downstream tasks. We utilize the spacy.Tokenizer.v1 configuration, which efficiently handles tokenization according to spaCy’s standards.

Next, we handle vector initialization. In this setup, vectors map tokens to high-dimensional representations, which helps capture semantic meaning during training. The data by converting the text and its annotations into a format compatible with spaCy.

3.3 Model Training

Our training utilizes the SpaCy pipeline configured with a transformer model and a transition-based parser for NER tasks. The deberta-v3-base model has been selected for the main transformer architecture, offering robust contextual embeddings for token-level classification (He et al., 2021).

Hyperparameters for the training are optimized based on performance on the development set. The key hyperparameters can be seen in Table 2.

3.4 Results and Discussion

We found that models utilizing spaCy achieved better results compared to those without it. Additionally, BERT base models outperformed other models in our experiments. However, we discov-

Model	F1-score
roberta-base	52.55
nlpaueb/legal-bert-base-uncased	53.29
lexlms/legal-roberta-base	54.80
lexlms/legal-roberta-base (Alibaba-NLP/gte-large-en-v1.5)	62.69
roberta-base with spacy	80.49
deberta-v3-base with spacy	86.37

Table 3: Comparison of F1-score in various models for the NER subtask

Model	F1-score
Nowj	0.416
Flawless Lawgic	0.402
<i>UOttawa</i>	<i>0.402</i>
Baseline	0.381
Masala-chai	0.380
UMLaw&TechLab	0.321
Bonafide	0.305

Table 4: Comparison of top 5 teams results on the hidden test set for the NER subtask, measured by F1-score (Hagag et al., 2024).

ered that initializing embeddings from Hugging Face leaderboard embeddings did not lead to improved results. Table 3 compares the F1-scores of various models on the labelled test data.

3.5 Direct Comparison to Related Work

The organizers of the shared task provided a hidden test set, on which our model achieved an F1-score of 0.402, securing third place in the competition. The performance of the top five teams is presented in Table 4.

4 Subtask B: Legal Natural Language Inference (L-NLI)

The goal of this subtask is to automatically classify the relationships between different legal texts. Specifically, we aim to determine whether a legal premise, such as a summary of a legal complaint, entails, contradicts, or remains neutral with respect to a given hypothesis, like an online review. This task, termed Legal Natural Language Inference (L-NLI), involves sentence-pair classification to assess these relationships. By creating an NLI corpus tailored for legal documents, we facilitate applications like legal case matching and automated legal reasoning. Detailed task definitions and datasets

are provided in (Bernsohn et al., 2024) and related resources.

4.1 Dataset Details

The LegalLensNLI dataset, provided by the organizers of the shared task, is specifically designed to explore the connections between legal cases and the individuals affected by them, with a particular focus on class action complaints. This dataset contains 312 entries. A comprehensive description of the dataset collection process can be found in (Bernsohn et al., 2024). For this subtask, only the training set is included, and the validation set is separated into four specific domains, as outlined in Table 6.

4.2 Preprocessing

This subtask has a different objective compared to Subtask A, so SpaCy may not perform well for this task. In this subtask, we began by loading the ynie/roberta-large-snli_mnli_fever_anli_R1_R2_R3-nli model (Nie et al., 2020) using the AutoTokenizer and AutoModel classes from the transformers library. The AutoTokenizer class was employed to tokenize the input sentences, converting them into a format suitable for the roberta-large model. The tokenization process involved splitting the text into tokens and converting them into numerical representations, which are then padded and truncated to a consistent length. This ensures that the input sequences are properly aligned when fed into the model.

Following tokenization, we implemented a method to encode the combined premise and hypothesis sentences for both the Roberta model and a CNN model. The CNN model required a different form of input preparation, where the combined texts were tokenized and encoded to maintain the sequence’s structure for CNN processing. These tokenized datasets were then converted into PyTorch tensors and mapped accordingly, enabling their use in a combined model that integrates both the Roberta model and the CNN. Subtask B involves finding the similarity between the hypothesis and premises. By using a CNN model to highlight the keywords in sentences, the combined model may perform better.

4.3 Model Training

Our combined model architecture integrates the ynie/roberta-large-snli_mnli_fever_anli_R1_R2_R3-nli model

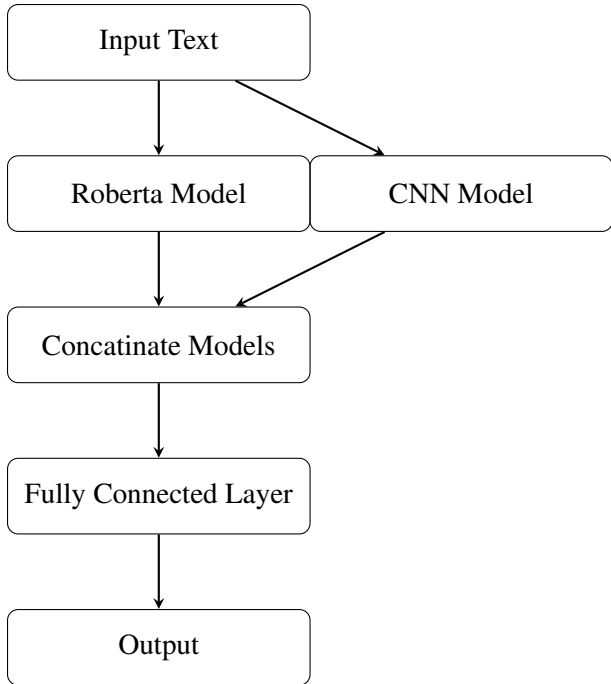


Figure 1: Diagram of combined model (Roberta and CNN)

with a custom-built CNN model for keyword detection. The Roberta model is responsible for capturing contextual information from the text, while the CNN model detects important keyword patterns within the input text. The outputs from both models are concatenated and passed through a fully connected layer (softmax) to produce the final classification decision, the architecture of the model can be seen in Figure 1. In more detail, the RoBERTa model consists of one embedding layer and 24 Transformer encoder layers, while the CNN model includes one embedding layer and three convolutional layers, each with a different filter size (2, 3, 4), followed by a fully connected layer. In total, we have 31 layers Training was conducted using the Trainer class from the transformers library, which facilitated the fine-tuning of the model. We defined specific hyperparameters, that can be seen Table 5. The model was evaluated at the end of each epoch, with the best model being saved based on the F1-score. The training process also included strategies for early stopping and warmup steps to optimize performance.

This approach combines the strengths of both the Roberta model and CNN, allowing for a more comprehensive analysis of the text data. The fine-tuning process ensures that the model is well-suited for the specific task of classifying legal text as 'Entailed', 'Neutral', or 'Contradict'.

Hyperparameter	Value
Learning Rate	2e-5
Batch Size (train and Eval)	4
Number of Epochs	20
Weight Decay:	0.01

Table 5: Hyperparameters of the fine-tuned model for subtask A (L-NLI)

Model	CP	Privacy	TCPA	Wage	Avg
Falcon 7B	87.2	84.5	83.9	68.5	81.02
without cnn	87.23	85.48	83.88	90.6	86.77
roberta-base	82.9	62.0	69.5	69.7	71.02
Our model	84.4	90	84	96	88.6

Table 6: Comparison of F1-score on the validation set for various models for the NLI task for specific-domain (Consumer Protection, Privacy,TCPA and Wage)

4.4 Results and Discussion

We found that pre-trained NLI models can perform significantly better than vanilla models and LLMs. Falcon 7B and RoBERTa base are the best-performing models for LLMs and vanilla models, respectively, as shown in Table 6. The validation set has been selected to be domain-specific, based on legal_act.

4.5 Direct Comparison to Related Work

The shared task organizers evaluated the models using a hidden test set, where our model attained an F1-score of 0.724, placing fifth in the competition. The results of the top five teams are detailed in Table 7.

5 Conclusion and Future Work

Our experiments demonstrated the success of transformer models, such as RoBERTa and DeBERTa, in handling complex legal tasks, including violation detection and inference. In Subtask A (L-NER), incorporating DeBERTa into the spaCy pipeline

Model	F1-score
1-800-Shared-Tasks	0.853
Baseline	0.807
Semantists	0.785
Nowj	0.746
UOttawa	0.724

Table 7: Performance of the leading 5 teams on the hidden test set in the NLI subtask, measured by F1-score (Hagag et al., 2024).

yielded strong results for legal named entity recognition. In Subtask B (L-NLI), combining RoBERTa with CNN for keyword detection boosted classification accuracy.

However, despite using robust models, generalizing to unseen cases proved challenging, particularly with nuanced legal language. While the CNN improved phrase detection, more advanced methods, like attention mechanisms, may further enhance performance.

Future work should explore architectures fine-tuned on legal texts or combine transformers with graph models to capture legal relationships. Additionally, leveraging LLMs like GPT-4 could improve legal reasoning.

References

- Hidelberg O Albuquerque, Ellen Souza, Adriano LI Oliveira, David Macêdo, Cleber Zanchettin, Douglas Vitório, Nádia FF da Silva, and André CPLF de Carvalho. 2023. On the assessment of deep learning models for named entity recognition of brazilian legal documents. In *EPIA Conference on Artificial Intelligence*, pages 93–104. Springer.
- Dor Bernsohn, Gil Semo, Yaron Vazana, Gila Hayat, Ben Hagag, Joel Niklaus, Rohit Saha, and Kyril Truskovskiy. 2024. [LegalLens: Leveraging LLMs for legal violation identification in unstructured text](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2129–2145, St. Julian’s, Malta. Association for Computational Linguistics.
- Luiz Henrique Bonifacio, Paulo Arantes Vilela, Gustavo Rocha Lobato, and Eraldo Rezende Fernandes. 2020. A study on the impact of intradomain finetuning of deep language models for legal named entity recognition in portuguese. In *Intelligent Systems: 9th Brazilian Conference, BRACIS 2020, Rio Grande, Brazil, October 20–23, 2020, Proceedings, Part I 9*, pages 648–662. Springer.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Can Çetindağ, Berkay Yazıcıoğlu, and Aykut Koç. 2023. Named-entity recognition in turkish legal texts. *Natural Language Engineering*, 29(3):615–642.
- Harshil Darji, Jelena Mitrović, and Michael Granitzer. 2023. German bert model for legal named entity recognition. *arXiv preprint arXiv:2303.05388*.
- Reto Gubelmann, Aikaterini-Lida Kalouli, Christina Niklaus, and Siegfried Handschuh. 2023. When truth matters-addressing pragmatic categories in natural language inference (nli) by large language models (llms). In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (* SEM 2023)*, pages 24–39.
- Ben Hagag, Liav Harpaz, Gil Semo, Dor Bernsohn, Rohit Saha, Pashootan Vaezipoor, Kyril Truskovskiy, and Gerasimos Spanakis. 2024. [Legallens shared task 2024: Legal violation identification in unstructured text](#). *Preprint*, arXiv:2410.12064.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). *Preprint*, arXiv:2111.09543.
- Mi-Young Kim, Juliano Rabelo, Housam Khalifa Bashier Babiker, Md Abed Rahman, and Randy Goebel. 2024. Legal information retrieval and entailment using transformer-based approaches. *The Review of Socionetwork Strategies*, 18(1):101–121.
- Shao-Man Lee, Yu-Hsiang Tan, and Han-Ting Yu. 2023. Learner: few-shot legal argument named entity recognition. In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, pages 422–426.
- Elena Leitner, Georg Rehm, and Julian Moreno-Schneider. 2019. Fine-grained named entity recognition in legal documents. In *International conference on semantic systems*, pages 272–287. Springer.
- Pedro Henrique Luz de Araujo, Teófilo E de Campos, Renato RR de Oliveira, Matheus Stauffer, Samuel Couto, and Paulo Bermejo. 2018. Lener-br: a dataset for named entity recognition in brazilian legal text. In *Computational Processing of the Portuguese Language: 13th International Conference, PROPOR 2018, Canela, Brazil, September 24–26, 2018, Proceedings 13*, pages 313–323. Springer.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. Adversarial NLI: A new benchmark for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Riccardo Pozzi, Riccardo Rubini, Christian Bernasconi, and Matteo Palmonari. 2023. Named entity recognition and linking for entity extraction from italian civil judgements. In *International Conference of the Italian Association for Artificial Intelligence*, pages 187–201. Springer.
- Yiu Kei Tang. 2023. Natural language inference transfer learning in a multi-task contract dataset: In the case of contractnli: a document information extraction system.
- Marco Valentino and André Freitas. 2024. On the nature of explanation: An epistemological-linguistic

perspective for explanation-based natural language inference. *Philosophy & Technology*, 37(3):88.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.

Xinrui Zhang, Xudong Luo, and Jiaye Wu. 2023. A roberta-globalpointer-based method for named entity recognition of legal documents. In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.