

# SemEval-2023 Task 6: LegalEval - Understanding Legal Texts

Ashutosh Modi<sup>¶</sup> Prathamesh Kalamkar<sup>†</sup> Saurabh Karn<sup>‡</sup>  
Aman Tiwari<sup>†</sup> Abhinav Joshi<sup>¶</sup> Sai Kiran Tanikella<sup>¶</sup>  
Shouvik Kumar Guha<sup>◊</sup> Sachin Malhan<sup>‡</sup> Vivek Raghavan<sup>\*</sup>

<sup>¶</sup>Indian Institute of Technology Kanpur (IIT-K),

<sup>†</sup>Thoughtworks, <sup>‡</sup>Agami, <sup>\*</sup>EkStep, <sup>◊</sup>NUJS, West Bengal

{ashutoshm, ajoshi, tskiran}@cse.iitk.ac.in,

{prathamk, aman.tiwari}@thoughtworks.com,

{saurabh, sachin}@agami.in,

shouvikkumarguha@nujs.edu, vivek@ekstep.org

## Abstract

In populous countries, pending legal cases have been growing exponentially. There is a need for developing NLP-based techniques for processing and automatically understanding legal documents. To promote research in the area of Legal NLP we organized the shared task LegalEval - Understanding Legal Texts at SemEval 2023. LegalEval task has three sub-tasks: Task-A (Rhetorical Roles Labeling) is about automatically structuring legal documents into semantically coherent units, Task-B (Legal Named Entity Recognition) deals with identifying relevant entities in a legal document and Task-C (Court Judgement Prediction with Explanation) explores the possibility of automatically predicting the outcome of a legal case along with providing an explanation for the prediction. In total 26 teams (approx. 100 participants spread across the world) submitted systems paper. In each of the sub-tasks, the proposed systems outperformed the baselines; however, there is a lot of scope for improvement. This paper describes the tasks, and analyzes techniques proposed by various teams.

## 1 Introduction

In populous countries (e.g., India), pending legal cases have grown exponentially. Due to the nature of the legal domain, it may not be possible to automate the entire judicial pipeline completely; nevertheless, many intermediate tasks can be automated to augment legal practitioners and hence expedite the system. For example, legal documents can be processed with the help of Natural Language Processing (NLP) techniques to organize and structure the data to be amenable to automatic search and retrieval. However, legal texts are different from commonly occurring texts typically used to train NLP models. Legal documents are quite long (tens and sometimes hundreds of pages), making it challenging to process (Malik et al., 2021). Another challenge with legal documents is that they

use different lexicons compared to standard texts. Moreover, legal documents are manually typed in countries like India and are highly unstructured and noisy (e.g., spelling and grammatical mistakes). Above mentioned challenges make it difficult to apply existing NLP models and techniques directly, which calls for the development of legal domain-specific techniques.

To promote research and development in the area of Legal-AI we proposed shared tasks via “LegalEval - Understanding Legal Texts”. This paper gives details about the shared task and our experience in organizing LegalEval. In total 26 teams submitted systems papers in the shared task. Though the documents pertained to the Indian legal system, teams across the globe participated showing an interest in developing legal-NLP techniques that could later be generalized to other legal systems as well.

The paper is organized as follows. In Section 2, we outline work done in the legal-NLP domain. Section 3 describes the tasks in detail. This is followed by details of the corpora (used for various sub-tasks) in Section 4. Details of models developed by each of the teams are discussed in Section 5. Section 6 analyzes the results furnished by various techniques for each of the sub-task. Finally, Section 7 summarizes and concludes the paper.

## 2 Related Work

In recent years, Legal NLP has become an active area for research. Researchers have worked on various research problems pertaining to the legal domain such as Legal Judgment Prediction (Zhong et al., 2020; Malik et al., 2021; Chalkidis et al., 2019; Aletras et al., 2016; Chen et al., 2019; Long et al., 2019; Xu et al., 2020; Yang et al., 2019a; Kapoor et al., 2022), Rhetorical Role Labeling (Bhattacharya et al., 2019; Saravanan et al., 2008; Malik et al., 2022a; Kalamkar et al., 2022b), Summarization (Tran et al., 2019), Prior Case Retrieval

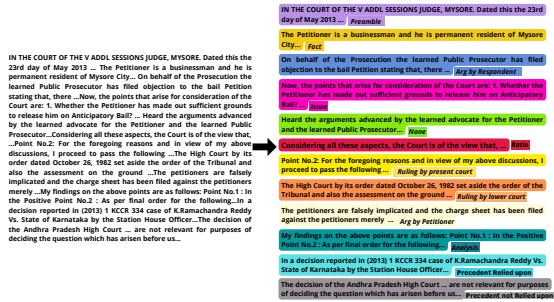


Figure 1: Example of document segmentation via Rhetorical Roles labels. On the left is an excerpt from a legal document, and on the right is a document segmented and labeled with rhetorical role labels.

and Statute Retrieval (Rabelo et al., 2022, 2020; Kim et al., 2019). Due to space limitations, we do not go into details of various tasks and allude to surveys on Legal-AI by Westermann et al. (2022); Villata et al. (2022); Colonna (2021). There have been some works in the area of Rhetorical Role Labeling and Legal Judgement Prediction; however, the research problems remain far from being solved, hence we organize shared tasks on these tasks. Moreover, Legal-NER (Named Entity Recognition) is an unexplored area, and we would like to draw the community’s attention to it.

### 3 Tasks Description

In LegalEval 2023, we proposed the following three shared tasks, which focus on automatically understanding legal documents. The documents were in English.

**Task A: Rhetorical Roles (RR) Prediction:** Given that legal documents are long and unstructured, we proposed a task for automatically segmenting legal judgment documents into semantically coherent text segments, and each such segment is assigned a label such as a preamble, fact, ratio, arguments, etc. These are referred to as Rhetorical Roles (RR). Concretely, we propose the task of Rhetorical Role Prediction: the task is to segment a given legal document by predicting (Figure 1) the rhetorical role label for each sentence (Malik et al., 2022b; Kalamkar et al., 2022b). We target 13 RRs as outlined in our previous work Kalamkar et al. (2022b): Preamble, Facts, Ruling by Lower Court, Issues, Argument by Petitioner, Argument by Respondent, Analysis, Statute, Precedent Relied, Precedent Not Relied, Ratio of the decision, Ruling by Present Court and None. For the definition of the RRs,

Entity	Description
Court	Name of the courts (supreme, high court, district, etc.)
Date	Dates mentioned in the. Judgement
Police Station	Name of Police Stations
Organization	Name of organization apart from court & police stations. E.g., private companies, banks, etc.
Statute	Act or law under which case is filed
Provision	Sections, articles or rules under the statute
Precedent	Higher Court cases referred
Case number	Number of the case
Petitioner Name	Names of Appellant/ Petitioner.
Respondent Name	Names of Respondent
Judge Name	Name of judge
Lawyer Name	Name of Lawyers
Witness Name	Name of witnesses
Others	Names of other people that don’t belong to any other categories above

Table 1: Legal Entity Types

please refer to Kalamkar et al. (2022b).

**RR Task Relevance:** The purpose of creating a rhetorical role corpus is to enable an automated understanding of legal documents by segmenting them into topically coherent units (Figure 1). This segmentation is a fundamental building block for many legal AI applications like judgment summarization, judgment outcome prediction, precedent search, etc.

**Task B: Legal Named Entity Recognition (L-NER):** Named Entity Recognition is a widely studied problem in Natural Language Processing (Li et al., 2018; Dozier et al., 2010). However, legal documents have peculiar entities like names of the petitioner, respondent, court, statute, provision, precedents, etc. These entity types are not recognized by the standard Named Entity Recognizer. Hence there is a need to develop a Legal NER system; consequently, we proposed Task B as Legal Named Entity Recognition (L-NER). In this task, we target 14 legal named entities as described in our previous work Kalamkar et al. (2022a) and also shown in Table 1. An example of the task is shown in Figure 2.

**Task Relevance:** Similar to the importance of NER in the general NLP domain, Legal NER is equally essential in the legal domain. For example, L-NER is the first step in extracting relevant entities for information extraction and retrieval-based tasks. Blackstone (<https://github.com/ICLRandD/Blackstone>) provides a li-

State of Kerala **Petitioner** ...Appellant  
 -versus-  
 Raneef **Respondent** ...Respondent  
 Judgment  
 Markandey Kalju, J. **Judge**

1. Leave granted.
2. Heard learned counsel for the parties.
3. The appellant has filed this appeal challenging the impugned order of the Kerala High Court. **Court** dated 1.9.2010 granting bail to the respondent, Dr. Raneef **Respondent** who is a medical practitioner (dentist) in Ernakulam district in Kerala, and is accused in crime no. 704 of 2010 **Case Number** of P.S. Muvattupuzha for 2 offences under various provisions of the I.P.C., the Explosive Substances Act **Statute**, and the Unlawful Activities (Prevention) Act. **Statute**

Figure 2: Example of Named Entities in a court judgment

brary for L-NER; however, it is limited to only UK cases and does not generalize well to case documents from other countries; moreover, the set of entities covered is quite limited.

**Task C: Court Judgment Prediction with Explanation (CJPE):** In recent years, in order to augment a judge in predicting the outcome of a case, the task of legal judgment prediction has been an active area of research (Malik et al., 2021). However, given the nature of the legal process, judgment prediction without providing an explanation for the judgments is not of much practical relevance to legal practitioners. We proposed the task of Legal Judgement Prediction with Explanation, where, given a legal judgment document, the task involves automatically predicting the outcome (binary: accepted or denied) of the case and also providing an explanation for the prediction. The explanations are in the form of relevant sentences in the document that contribute to the decision. This task is an extension of one of our previous works (Malik et al., 2021). **We divided the task into two sub-tasks: (1) Court Judgment Prediction and (2) Explanations for the prediction.** As done in our previous work, we take all the steps to remove any bias in the dataset and address the ethical concerns associated with the task.

*Task Relevance:* The volume of legal cases has been growing exponentially in many countries; if there is a system to augment judges in the process, it would be of immense help and expedite the process. Moreover, if the system provides an explanation for the decision, it would help a judge make an informed decision about the final outcome of the case.

### 3.1 Evaluation

The rhetorical roles task (a multiclass prediction problem) is evaluated using a weighted F1

score based on the test data. For the CJPE task, the evaluation for judgment prediction (binary classification) is done using the standard F1 score metric, and for the explanation sub-task, we use ROUGE-2 score (Malik et al., 2021) for evaluating the machine explanations with respect to the gold annotations. For L-NER, we use standard F1 score metrics (Segura-Bedmar et al., 2013). Automatic evaluations of the submitted models are done using CodaLab, and a leaderboard is maintained (<https://codalab.lisn.upsaclay.fr/competitions/9558>).

## 4 Corpora

We focus on freely and publicly available Indian legal documents in English. The documents are obtained from Indian Legal search portal IndianKanoon (<https://indiankanoon.org/>). India follows a common-law system, which is similar to the one in the U.S. and the U.K. The tasks organizers with the help of legal experts in the past, have created the annotated datasets for RR (Kalamkar et al., 2022b; Malik et al., 2022b), L-NER (Kalamkar et al., 2022a) and CJPE (Malik et al., 2021). The data is annotated by senior law students, professors, and legal practitioners, thus ensuring the high quality of the data. We annotated around 250 judgments which are likely to have 850k tokens and roughly 10k entities. The RR data has about 265 legal documents with 841,648 tokens and 26,304 sentences, with an average of 3176 tokens per document. It is split into the train, validation, and test sets with an approximate 70-10-20 split. L-NER corpus has about 14,444 judgment documents annotated with 17,485 named entities belonging to 14 categories. The CJPE corpus has about 34K documents, and these are split into the train, validation, and test sets.

## 5 Participating Systems

We organized the shared task using the CodaLab platform. In total, 40 teams participated in various sub-tasks and finally 26 teams (approx. 100 participants) submitted systems description papers. Overall, there were 17 submissions for Task-A (RR) and 11 submissions each for Task-B (NER) and Task-C (CJPE). Teams other than from India also participated in the shared task. The demographic distribution of teams is shown in Figure 3. Next, we summarize the proposed approaches by various teams for each of the tasks.




Figure 3: Demographics of the participating systems

### 5.1 Task-A (Rhetorical Role Prediction)

In total, seventeen teams provided a description of their approach. We observed a wide variety of approaches for Task-A, like BiLSTM-based, Transformer-based, and GNN-based architectures, out of which BiLSTM was one of the popular choices (6 teams). Another common trend was the use of pre-trained transformer-based architectures as a backbone, e.g., BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), LegalBERT (Chalkidis et al., 2020), InLegalBERT (Paul et al., 2022).

AntContentTech (Huo et al., 2023) proposed a system based on Bi-LSTM with CRF layers and used domain-adaptive pre-training with auxiliary-task learning with pre-trained LegalBERT as a backbone. They also experimented with augmentation strategies like shuffling precedents among the documents and back translation. This helps to improve performance.

VTCC-NLP (Nguyen et al., 2023) followed a unique strategy of using Graph Neural Nets with sentence embeddings. Moreover, the proposed method used multitask learning with label shift prediction and contrastive learning to separate close labels like *ARG\_petitioner* and *ARG\_respondent*. Combining entity-based GNN (GAT (Veličković et al., 2018)) with multitask learning and contrastive learning helped to improve performance.

TeamShakespeare (Jin and Wang, 2023) highlighted the significance of considering inter and intra-sentence representations in RR predictions. The method also reported improvements achieved by replacing BERTbase with LegalBERT as a backbone in the architecture.

NLP-Titan (Kataria and Gupta, 2023) followed a two-stage Hierarchical strategy where in the first stage, similar classes were merged, and a sequence sentence classification algorithm

based on Hierarchical Sequential Labeling Network (HSLN) was used. In the second stage, the SetFit model is trained over the fine-grained merged classes.

YNU-HPCC (Chen et al., 2023) proposed the use of LegalBERT with attention pooling and a CRF module. Predictions are made using the Deep ensembles (Lakshminarayanan et al., 2017) strategy, where multiple models were trained with different seeds followed by a voting scheme.

TeamUnibo (Noviello et al., 2023) used a pre-trained BERT with two BiLSTM layers followed by average pooling and a linear layer. The proposed method was motivated to maintain topical coherence by taking the surrounding sentences as context for RR classification.

IRIT\_IRIS\_(A) (Lima et al., 2023) proposed a unique style of segmenting long documents referred to as DFCSC (Dynamic Filled Contextualized Sentence Chunks), where every chunk contained a variable number of core and edge sentences. The paper further described 3 variants of the proposed method and reported the results.

LRL\_NC (Tandon and Chatterjee, 2023) used a hierarchical sequence encoder with contextual sentence representations followed by a CRF layer for classifying sentences into RR categories.

ResearchTeam\_HCN (Ningthoujam et al., 2023) used back translations to balance the imbalanced classes and proposed using attention pooling with sentence transformers.

UOttawa.NLP23 (Almuslim et al., 2023) highlighted improvements with a Hierarchical BiLSTM-CRF when compared to other transformer-based models.

NITK\_Legal (Sindhu et al., 2023) also used a BiLSTM architecture with the CRF layer, which follows a hierarchical structure. The paper also reported its findings on various sentence encoding approaches like BiLSTM embeddings, sentence2vec (sent2vec) (Moghadasi and Zhuang, 2020), Universal Sentence Encoder (Cer et al., 2018), Sentence BERT (SBERT) (Reimers and Gurevych, 2019), trigram embeddings, weighted trigram and handcrafted features + sent2vec and reported the importance of handcrafted features along with sent2vec embeddings for RR classification.

TüeReuth-Legal (Manegold and Girrbach, 2023) explored the learning of a continuous distribution over positions in the document to indicate the likelihood of respective RRs appearing at cer-

tain positions. The method further incorporated domain knowledge through hard-coded lexical rules to improve performance.

NITS<sub>Legal</sub> (Jain et al., 2023) created a different version of the dataset for capturing local context and used a BiLSTM-based sentence classification approach, highlighting the significance of context in RR prediction.

Steno AI (Gupta et al., 2023) created a graphical structure for capturing the relationship between the sentences in the document and applied GNN-based classification to predict the rhetorical roles.

CSECU-DSG (Tasneem et al., 2023) framed the task as a multi-class sequence classification task. In the proposed approach the authors utilized the fine-tuned Legal-BERT for classifying legal sentences into 13 rhetorical roles.

LTRC (Baswani et al., 2023) used several algorithms to classify sentence embeddings into different classes of Rhetorical Roles. The authors showed that using 512 instead of 256 as embedding size gives better results. Further, they showed that enriching the embeddings and adding a CRF layer to the classifier helps to improve results for classes with fewer samples.

UO-LouTAL (Bosch et al., 2023) proposed a Bag-of-Words-based system and used NER labels to replace entities with their types to reduce sentence variability.

## 5.2 Task-B (Named Entity Recognition)

For the NER task, out of 17 participating teams, 11 teams submitted system description papers. Most teams (7) used variations of the BERT-based transformers like RoBERTa, DeBERTaV3 (He et al., 2021), LegalBERT, Albert (Lan et al., 2019). techniques based on XLNet (Yang et al., 2019b), LUKE-base, and Legal-LUKE are also implemented. BiLSTM and CRF layers are used in the final layers on the top of transformers.

ResearchTeam\_HCN (Ningthoujam et al., 2023) proposed two baseline models and one main model with an MTL (Multi-Task Learning) framework. They used data augmentation (back-translation) as a way to make up for the imbalance in class data. Further, they used a general-purpose NER system to aid the Legal-NER model. For example, using rule-based techniques they assigned labels to an entity predicted by general-purpose NER.

VTCC-NER (Tran and Doan, 2023) proposed

the use of ensembles to improve the NER predictions via two approaches. The first approach considered the use of multiple transformer backbones (RoBERTa, BERT, Legal-BERT, InLegalBERT, and XLM-RoBERTa) in the spaCy framework to make the predictions, which are further combined using a weighted voting scheme to generate the NER predictions. In the second approach, they adapted MultiCoNER (Multilingual Complex Named Entity Recognition (Malmasi et al., 2022)) with multiple pre-trained backbones to generate the contextualized features, which are further concatenated with dependency parsing features and part-of-speech features. The obtained concatenated features were further passed through a CRF module to generate the predictions.

Jus Mundi (Cabrera-Diego and Gheewala, 2023) proposed the method FEDA (Frustratingly Easy Domain Adaptation), and trained it over multiple legal corpora. The FEDA layer consisted of a GeLU activation layer, followed by a linear layer and layer normalization. The method extracted features from DeBERTa V3 and passed them through two splits: the general FEDA layer and the specialized FEDA layer. The specialized FEDA layer also incorporated features from the general FEDA layer, and the final output was passed through a CRF module to generate the predictions.

UOttawa.NLP23 (Almuslim et al., 2023) converted the text to B-I-O format and tokenized it using the Legal-Bert (Chalkidis et al., 2020). The authors proposed a token classification approach where each token was classified into a particular entity type in the dataset.

TeamUnibo (Noviello et al., 2023) presented an approach to handle raw text without any pre-processing. A sliding window approach handled long texts and avoided input truncation. The system used the B-I-O labeling format and mapped labels from words to tokens. Several baseline models, including RoBERTa, InLegalBERT, and XLNet, were tested, along with two custom models that use RoBERTa or XLNet as the transformer backbone, followed by a BiLSTM and CRF layers. The best-performing model was found to be RoBERTa - BiLSTM - CRF and XLNet - BiLSTM - CRF.

AntContentTech (Huo et al., 2023) used a SciBERT-HSLN to get word embeddings of each sentence which were then passed through a BiLSTM layer and an attention-based pooling layer to get sentence representations. These representations

are directly fed into a CRF layer to perform the span classification.

PoliToHFI (Benedetto et al., 2023) applied the entity-aware attention mechanism implemented in the LUKE model (Yamada et al., 2020).

Ginn-Khamov (Ginn and Khamov, 2023) proposed a three step augmented approach. First, a classification model, where the sentence class was added as an input token, then a twin transformer model with two identical NER models was used and each batch of sentences was split into preamble and judgment using the previous classification model. The sentence class was used to pick one of the two sub-models, and finally, a modified token classification linear layer took the sentence class provided by the classifier predictions to learn the differences between documents.

TeamShakespeare (Jin and Wang, 2023) introduced Legal-Luke, a legal entity recognition model based on the bidirectional transformer encoder of LUKE (Yamada et al., 2020). The model generated the Legal contextual embeddings for words and entities by adding the token embeddings, type embeddings. Position embeddings were added to compute the word and entity embeddings.

UO-LouTAL (Bosch et al., 2023) aimed to achieve interpretability in the L-NER task by using a standard CRF. The input to the CRF model was a structure containing token features retrieved from the text using SpaCy, and the features of its preceding and succeeding tokens. Optionally, post-processing was applied to extract specific patterns such as dates or case numbers.

Legal\_try (Zhao et al., 2023) proposed the use of multiple heterogeneous models, which were further selected using a recall score. Later, all the selected models were used in a voting scheme to generate the predictions. The heterogeneous models included ALBERT – BiLSTM – CRF, ALBERT – IDCNN – CRF, and RoBERTa (dynamic fusion) – BiLSTM/IDCNN – CRF.

Nonet (Nigam et al., 2023) used a combination of results produced by the custom-trained spaCy model and fine-tuned BERT-CRF model for the final prediction. spaCy was trained by replacing the RoBERTa embeddings with the law2vec embeddings and the LegalBERT was used in the pipeline for fine-tuning the model. The BERT-CRF model was trained over the pre-processed text using the BIO format. If the span of the entities predicted by both models intersected then the entity with a larger

span was considered. For the rest, a union of the result produced by both models was considered.

### 5.3 Task C1 and C2: Legal Judgment Prediction and Legal Judgement Prediction with Explanation

For the Legal Judgement Prediction task (C1) transformers were used as the base for 7 out of 8 participating systems. Longformer, RoBERTa, InLegalBERT, LegalBERT, and CaseLawBERT are used. One team `nclu_team` proposed the use of CNN to solve this task. In the Explanation for Judgement Prediction task (C2), various techniques were employed to generate the explanations. Some of those included the use of occlusion, clustering of similar sentences and picking sentences from the clusters, using relevance scores based on the proportion of legal entities, use of attention scores, extractive summarization, span lengths, and generation using the T5-based transformer model.

Viettel-AI (Hoang et al., 2023) used LongFormer (Beltagy et al., 2020) to process documents. High and low-frequency words were filtered out and a TF-IDF representation of the document was created. The TF-IDF vector representation was used along with LongFormer to learn representations in an end-to-end fashion. The authors further extended the proposed method for explanations by using a trained judgment prediction model and masking the sentences present in the document. With the help of this occlusion technique, the authors ranked the sentences based on the difference in the judgment prediction probability scores.

IRIT\_IRIS\_(C) (Prasad et al., 2023) proposed a multi-level encoder-based classifier using a pretrained BERT-based encoder as a backbone and fine-tuned over the chunked version of the documents. Further, the extracted chunk embeddings were used to represent the entire document, followed by another transformer and recurrent layer in level-3 to learn the relation between the chunk representations. For explanations, the proposed method followed the idea of occlusion sensitivity as an extension to the judgment prediction module.

UOttawa.NLP23 (Almuslim et al., 2023) explored the performance of different transformer-based models and hierarchical BiLSTM-CRF models for legal text classification on the Indian Legal Document Corpus (ILDC). Different models, including Legal-BERT, CaseLawBERT, and In-CaseLawBERT were fine-tuned and used as clas-

Rank	Team Name	Paper Submission	Model	Score
1	AntContentTech	yes	LegalBERT+BiLSTM+CRF	0.8593
2	xixilu556	no		0.8581
3	HJ	no		0.8497
4	VTCC-NLP	yes	InLegalBERT + contrastive learning + entity-based GNN	0.8389
5	TeamShakespeare	yes	Legal + BERT + HSLN	0.8343
6	NLP-Titan	yes	HSLN + BERT + CRF	0.8309
7	YNU-HPCC	yes	Legal-BERT + attention pool + CRF	0.8146
8	TeamUnibo	yes	InLegalBERT + BiLSTM	0.8112
9	IRIT_IRIS_(A)	yes	DFCSC ( Dynamic Filled Contextualized Sentence Chunks)	0.8076
10	peanut	no		0.8038
11	DeepAI	no		0.8028
12	LRL_NC	yes	RoBERTa + HDP + BiLSTM + CRF	0.7980
13	NU	no		0.7951
14	ResearchTeam_HCN	yes	Sentence Transformers + BiLSTM	0.765
15	UOttawa.NLP23	yes	Hierarchical BiLSTM-CRF	0.7465
16	nitk_legal	yes	Handcrafted Features + sent2vec	0.7400
17	TueReuth-Legal	yes	RoBERTa-Base / Legal- BERT + temporal distribution of labels	0.7335
18	NITS_Legal	yes	LegalBERT+BiLSTM	0.7143
19	Steno AI	yes	Context-based LegalBERT (best results)	0.7102
20	scholarly360	no		0.6902
21	aashsach	no		0.6602
22	CSECU-DSG	yes	LegalBERT	0.6465
23	LTRC	yes	LegalBERT+LDA / InCaseLawBERT	0.6133
24	ccidragons	no		0.5813
25	UO-LouTAL	yes	TF-IDF + SBERT	0.5681
26	achouhan	no		0.3506
27	viniciusvidal	no		0.2609

Table 2: Leaderboard Results for Task-A: Rhetorical Roles (RR) Prediction

sifiers. The best-performing model was CaseLawBERT + BiGRU. The successful performance of CaseLawBERT over InCaseLawBERT can be attributed to the pretraining corpus used. However, the models struggled to generalize to longer texts, and the validation results were better than the test results due to the longer texts in the test set. The authors used a masking technique with CaseLawBERT and BiGRU models to extract explanations for legal outcomes. They clustered similar sentences and picked one or two sentences from each cluster to form explanations using two different approaches. The authors concluded that their low score was due to selecting only one/two sentences per chunk, which resulted in the lower overlap.

PoliToHFI (Benedetto et al., 2023) used a three-staged approach where in the first stage, 4 transformer based encoders were employed to find the section exerting the most significant influence on the classification performance, and the best-performing encoder was identified. In the second stage, document representation was obtained by average pooling the sentence representations obtained using the chosen model. This representation was passed to the linear layers to compute the final prediction. Finally, an ensemble prediction for the test set was generated from the combination

of hierarchical models. They leveraged post hoc input attribution explanation methods to generate prediction explanations. The explanation was derived by adopting the post hoc attribution methods that leverage the input document and the model. Further, a relevance score based on each sentence’s proportion of legal entities was assigned to boost the explanations.

ResearchTeam\_HCN (Ningthoujam et al., 2023) followed a hierarchical model approach in which a Legal pre-trained RoBERTa model was fine-tuned using the last 512 tokens from each document. Later, the features were extracted using the fine-tuned Legal-RoBERTa model and were passed through the BiGRU model and a classification layer. For explanation, they generated relevant sentences using the T5-based transformer model and considered the task as a SEQ2SEQ approach.

AntContentTech (Huo et al., 2023) used a self-attention layer to generate contextualized sentence representations. A global token pre-pended to the input sequence was utilized to generate the final prediction of the judgment. The attention scores were used to extract the explanations. Based on these attention scores provided by the first token, the top 30% of sentences were selected.

nclu\_team (Rusnachenko et al., 2023) pro-

Rank	Team Name	Paper Submission	Model	Score
1	ResearchTeam_HCN	yes	RoBERTa-base	0.9120
2	VTCC-NER	yes	RoBERTa + Transition based Parser with dynamic weights	0.9099
3	DeepAI	no		0.9099
4	Jus Mundi	yes	DeBERTa V3 + FEDA + CRF	0.9007
5	Autohome AI	no		0.8833
6	uottawa.nlp23	yes	LegalBERT	0.8794
7	TeamUnibo	yes	XLNet + BiLSTM + CRF	0.8743
8	DamoAI	no		0.8627
9	AntContentTech	yes	LegalBERT + CRF ensemble	0.8622
10	xixilu556	no		0.8611
11	PoliToHFI	yes	LUKE-base / BERT-base	0.8321
12	Ginn-Khamov	yes	RoBERTa + Sentence class token	0.7265
13	TeamShakespeare	yes	legal- LUKE	0.6670
14	UO-LouTAL	yes	SpaCy (preprocessing) + CRF	0.6489
15	Nonet	yes	custom trained spaCy + BERT-CRF	0.5532
16	Legal_try	yes	ALBERT / RoBERTa + BiLSTM / IDCNN +CRF	0.5173
17	shihanmax	no		0.0186

Table 3: Leaderboard Results for Task B: Legal Named Entity Recognition (L-NER)

posed solving judgment prediction via Convolution Neural Network and used attention layers over CNN to generate the explanations. The proposed use of a sliding window over the most attentive part is unique in terms of using the ranking order of the obtained attention scores and it provided better and comparable results when compared to other explanation approaches.

Nonet (Nigam et al., 2023) used transformers like XLNet, Roberta, LegalBERT, InLawBert, InCaseLaw, and BERT-Large and also tried Hierarchical Transformer model architecture with a moving window approach to reduce the document into chunks. The CLS representation of these chunks was then used as input to sequential models (BiGRU + attention). In the Explanation for the prediction task, various keyword-based matching techniques along with Rhetorical Roles (ratio) were used to find the court judgment decision. Different span lengths from the ending portions of the legal judgments were considered for the explanation.

## 6 Results and Discussion

**Baselines:** We used transformer-based models as a baseline for each of the tasks. The baseline model for RR tasks was SciBERT-HSLN (Brack et al., 2021) (<https://github.com/Legal-NLP-EkStep/rhetorical-role-baseline>). We used RoBERTa base with the transition-based parser as the baseline for the NER task ([https://github.com/Legal-NLP-EkStep/legal\\_NER](https://github.com/Legal-NLP-EkStep/legal_NER)). Similarly, the baseline model for CJPE was a

Hierarchical model with XL-Net at the first level and Bi-GRU at the next level <https://github.com/Exploration-Lab/CJPE>.

**Task A:** Results for Task A are shown in Table 2. Overall, the use of BiLSTM layers with a CRF module seemed to be a popular, showing promising results. In general, the strategy of augmenting the dataset seemed to be a better choice which showed promising improvements and generalized predictions, and worked well on the unseen test set. One of the teams (VTCC-NLP) also used GNNs to capture the relationship between the entities to classify sentences into RR categories. The best-obtained results on the test set were around 86% (Team AntContentTech) followed by scores ranging from 84% to 80% where the majority of the methods used pre-trained transformer-based architectures trained on Legal text (e.g., LegalBERT, InLegalBERT) as the primary backbone for legal text representations. The baseline model has a performance of 79% F1 score. The proposed methods show an improvement over the baseline, however, there is further scope for improvement.

**Task B:** Results for Task B are shown in Table 3. Most of the teams used transformer-based backbone (e.g., RoBERTa, LegalBERT). Surprisingly, the use of the Legal-BERT-based performs similarly to standard pre-trained RoBERTa, highlighting the generic nature of named entity recognition tasks that may not require specialized legal knowledge. Augmentation techniques like back-translation seem to improve the performance of NER systems on Legal text. The best-performing system uses ensembles to further improve the re-



Rank	Team Name	Paper Submission	Model	Score
1	Viettel-AI	yes	LongFormer + TF-IDF	0.7485
2	IRIT_IRIS_(C)	yes	InLegalBERT+TransformerEncoder	0.7228
3	uottawa.nlp23	yes	CaseLawBERT + BiGRU	0.6782
4	PoliToHFI	yes	Transformer-based encoders + Avg. pool on segmented sentences	0.6771
5	maodou	no		0.6735
6	ResearchTeam_HCN	yes	RoBERTa + BiGRU	0.6671
7	AntContentTech	yes	LegalBERT	0.6416
8	nclu_team	yes	CNN + Attention + Sliding Window	0.6351
9	DamoAI	no		0.6326
10	xixilu556	no		0.6325
11	Nonet	yes	BERT / Hierarchical Transformer + BiGRU+Attention	0.5287

Table 4: Leaderboard Results for Task-C-1: Court Judgment Prediction

Rank	Team Name	Paper Submission	Model	Score
1	Nonet	yes	Keyword-based matching techniques + RR	0.5417
2	Viettel-AI	yes	Longformer + TF-IDF	0.4797
3	nclu_team	yes	CNN + Attention + Sliding Window	0.4789
4	xixilu556	no		0.4781
5	DamoAI	no		0.4781
6	uottawa.nlp23	yes	CaseLawBERT + BiGRU	0.4781
7	PoliToHFI	yes	post-hoc input attribution	0.4600
8	ResearchTeam_HCN	yes	T5 (modelled as a SEQ2SEQ approach)	0.4583
9	AntContentTech	yes	LegalBERT	0.4525
10	IRIT_IRIS_(C)	yes	InLegalBERT+TransformerEncoder	0.4000
11	maodou	no		0.3600

Table 5: Leaderboard Results for Task-C-2: Explanations for the Court Judgment Prediction

sults of the proposed system. In comparison to the baseline (F1 score of 91.1%), there was not much improvement in the best-performing system. This highlights the development of better techniques to improve the Legal-NER system further.

**Task C:** Table 4 and Table 5 show the results for the task of Legal Judgement Prediction (LJP) and Legal Judgment Prediction with Explanation (LJPE), respectively. For the LJP task, all the systems used transformer-based models, with the best-performing model using LongFormer, which was able to capture long-range dependencies in legal documents. For the LJPE task, systems used occlusion-based and attention-based methods for the explanation. The best-performing system used Rhetorical Roles to extract the ratio of the case, and that was subsequently used to provide the explanation. This is similar to the approach proposed by Malik et al. (2022a). Viettel-AI (Hoang et al., 2023) group achieved the highest F1 score of 74.85% in task C1 and Nonet (Nigam et al., 2023) has the highest score of 54.18% in task C2.

## 7 Conclusion

In this paper, we present the findings of ‘SemEval 2023 Task 6: LegalEval - Understanding Legal

Texts’, consisting of three Subtasks: i) Subtask A, Rhetorical Role identification, aimed to segment a legal document by assigning a rhetorical role to each of the sentence present in the document ii) Subtask B, Named Entity Recognition, for identifying NER tags present in a legal text and iii) Subtask C, Legal Judgment Prediction with Explanation, which aims to predict the final judgment of a legal case and in addition, provides explanations for the prediction. In total, 40 teams participated and 26 teams (approximately 100 participants) submitted systems papers and presented various models for each task. We provide an overview of methods proposed by the teams. Analysis of the systems and results points out more scope for improvement. We hope these findings will provide valuable insights to improve approaches for each of the sub-tasks.

## Acknowledgements

We would like to thank legal experts who contributed in annotating datasets. In particular, we thank Shreyangshi Gupta, Priyanka Choudhary, Srinjoy Das and Khushi Jain for annotating CJPE test set.

## References

- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lampos. 2016. Predicting judicial decisions of the european court of human rights: a natural language processing perspective. *PeerJ Comput. Sci.*, 2:e93.
- Intisar Almuslim, Sean Stilwell, Surya Kiran Suresh, and Diana Inkpen. 2023. uOttawa at SemEval-2023 Task 6: Deep Learning for Legal Text Understanding. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Pavan Baswani, Hiranmai Sri Adibhatla, and Manish Shrivastava. 2023. LTRC at SemEval-2023 Task 6: Experiments with Ensemble Embeddings. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Irene Benedetto, Alkis Koudounas, Lorenzo Vaiani, Eliana Pastor, Elena Baralis, Luca Cagliero, and Francesco Tarasconi. 2023. PoliToHFI at SemEval-2023 Task 6: Leveraging Entity-Aware and Hierarchical Transformers For Legal Entity Recognition and Court Judgment Prediction. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1401–1411, Toronto, Canada. Association for Computational Linguistics.
- Paheli Bhattacharya, Shounak Paul, Kripabandhu Ghosh, Saptarshi Ghosh, and Adam Wyner. 2019. Identification of rhetorical roles of sentences in indian legal judgments. *CoRR*, abs/1911.05405.
- Sébastien Bosch, Louis Estève, Joanne Loo, and Anne-Lyse Minard. 2023. UO-LouTAL at SemEval-2023 Task 6: Lightweight Systems for Legal Processing. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Arthur Brack, Anett Hoppe, Pascal Buschermöhle, and Ralph Ewerth. 2021. Sequential sentence classification in research papers using cross-domain multi-task learning. *arXiv e-prints*, pages arXiv–2102.
- Luis Adrián Cabrera-Diego and Akshita Gheewala. 2023. Jus Mundi at SemEval-2023 Task 6: Using a Frustratingly Easy Domain Adaption for a Legal Named Entity Recognition System. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural Legal Judgment Prediction in English. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. LEGAL-BERT: The muppets straight out of law school. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online. Association for Computational Linguistics.
- Huajie Chen, Deng Cai, Wei Dai, Zehui Dai, and Yadong Ding. 2019. Charge-based prison term prediction with deep gating network. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6362–6367, Hong Kong, China. Association for Computational Linguistics.
- Yu Chen, You Zhang, Jin Wang, and Xuejie Zhang. 2023. YNU-HPCC at SemEval-2023 Task 6: LEGAL-BERT Based Hierarchical BiLSTM with CRF for Rhetorical Roles Prediction. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Liane Colonna. 2021. Reflections on the use of ai in the legal domain. *Law and Business*, 1(1):1–10.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Christopher Dozier, Ravikumar Kondadadi, Marc Light, Arun Vachher, Sriharsha Veeramachaneni, and Ramdev Wudali. 2010. Named entity recognition and resolution in legal text. In *Semantic Processing of Legal Texts*, pages 27–43. Springer.
- Michael Ginn and Roman Khamov. 2023. Ginn-Khamov at SemEval-2023 Task 6, Subtask B: Legal Named Entities Extraction for Heterogenous Documents. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 420–425, Toronto, Canada. Association for Computational Linguistics.
- Anshika Gupta, Shaz Furniturewala, Vijay Kumari, and Yashvardhan Sharma. 2023. Steno AI at SemEval-2023 Task 6: Rhetorical Role Labelling of Legal Documents using Transformers and Graph Neural Networks. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1858–1862,

- Toronto, Canada. Association for Computational Linguistics.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- Thanh Dat Hoang, Chi Minh Bui, and Nam Khac-Hoai Bui. 2023. **Viettel-AI at SemEval-2023 Task 6: Legal Document Understanding with Longformer for Court Judgment Prediction with Explanation**. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 862–868, Toronto, Canada. Association for Computational Linguistics.
- Jingjing Huo, Kezun Zhang, Zhengyong Liu, Xuan Lin, Wenqiang Xu, Maozong Zheng, Zhaoguo Wang, and Song Li. 2023. **AntContentTech at SemEval-2023 Task 6: Domain-adaptive Pretraining and Auxiliary-task Learning for Understanding Indian Legal Texts**. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 402–408, Toronto, Canada. Association for Computational Linguistics.
- Deepali Jain, Malaya Dutta Borah, and Anupam Biswas. 2023. NITS\_Legal at SemEval-2023 Task 6: Rhetorical Roles Prediction of Indian Legal Documents via Sentence Sequence Labeling Approach. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Xin Jin and Yuchen Wang. 2023. TeamShakespeare at SemEval-2023 Task 6: Understand Legal Documents with Contextualized Large Language Models. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Prathamesh Kalamkar, Astha Agarwal, Aman Tiwari, Smita Gupta, Saurabh Karn, and Vivek Raghavan. 2022a. **Named Entity Recognition in Indian court judgments**. In *Proceedings of the Natural Language Processing Workshop 2022*, pages 184–193, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Prathamesh Kalamkar, Aman Tiwari, Astha Agarwal, Saurabh Karn, Smita Gupta, Vivek Raghavan, and Ashutosh Modi. 2022b. **Corpus for Automatic Structuring of Legal Documents**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4420–4429, Marseille, France. European Language Resources Association.
- Arnav Kapoor, Mudit Dhawan, Anmol Goel, Arjun T H, Akshala Bhatnagar, Vibhu Agrawal, Amul Agrawal, Arnab Bhattacharya, Ponnurangam Kumaraguru, and Ashutosh Modi. 2022. **HLDC: Hindi legal documents corpus**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3521–3536, Dublin, Ireland. Association for Computational Linguistics.
- Harsh Kataria and Ambuje Gupta. 2023. **NLP-Titan at SemEval-2023 Task 6: Identification of Rhetorical Roles Using Sequential Sentence Classification**. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1365–1370, Toronto, Canada. Association for Computational Linguistics.
- Mi-Young Kim, Juliano Rabelo, and Randy Goebel. 2019. Statute law information retrieval and entailment. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 283–289.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. **Simple and scalable predictive uncertainty estimation using deep ensembles**. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. 2018. **A survey on deep learning for named entity recognition**.
- Alexandre Lima, Jose G. Moreno, and Eduardo H. da S. Aranha. 2023. IRIT\_IRIS\_A at SemEval-2023 Task 6: Legal Rhetorical Role Labeling Supported by Dynamic-filled Contextualized Sentence Chunks. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized BERT pretraining approach**. *CoRR*, abs/1907.11692.
- Shangbang Long, Cunchao Tu, Zhiyuan Liu, and Maosong Sun. 2019. Automatic judgment prediction via legal reading comprehension. In *Chinese Computational Linguistics - 18th China National Conference, CCL 2019, Kunming, China, October 18-20, 2019, Proceedings*, volume 11856 of *Lecture Notes in Computer Science*, pages 558–572. Springer.
- Vijit Malik, Rishabh Sanjay, Shouvik Kumar Guha, Angshuman Hazarika, Shubham Nigam, Arnab Bhattacharya, and Ashutosh Modi. 2022a. **Semantic Segmentation of Legal Documents via Rhetorical Roles**. In *Proceedings of the Natural Language Processing Workshop 2022*, pages 153–171, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Vijit Malik, Rishabh Sanjay, Shouvik Kumar Guha, Angshuman Hazarika, Shubham Nigam, Arnab Bhattacharya, and Ashutosh Modi. 2022b. **Semantic segmentation of legal documents via rhetorical roles**. In *Proceedings of the Natural Language Processing Workshop 2022*, pages 153–171, Abu Dhabi,

- United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, and Ashutosh Modi. 2021. [ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4046–4062, Online. Association for Computational Linguistics.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022. Semeval-2022 task 11: Multilingual complex named entity recognition (multiconer). In *Proceedings of the 16th international workshop on semantic evaluation (SemEval-2022)*, pages 1412–1437.
- Henrik Manegold and Leander Gırrbach. 2023. [TüReuth Legal at SemEval-2023 Task 6: Modelling Local and Global Structure of Judgements for Rhetorical Role Prediction](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Mahdi Naser Moghadasi and Yu Zhuang. 2020. [Sent2vec: A new sentence embedding representation with sentimental semantic](#). In *2020 IEEE International Conference on Big Data (Big Data)*, pages 4672–4680. IEEE.
- Hiep Huu Nguyen, Hoang Ngo, and Nam Bui. 2023. [VTCC-NLP at SemEval-2023 Task 6: Long-Text Representation Based on Graph Neural Network for Rhetorical Roles Prediction](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1121–1126, Toronto, Canada. Association for Computational Linguistics.
- Shubham Kumar Nigam, Aniket Deroy, Noel Shallum, Ayush Kumar Mishra, Anup Roy, Shubham Kumar Mishra, Arnab Bhattacharya, Saptarshi Ghosh, and Kripabandhu Ghosh. 2023. [Nonet at SemEval-2023 Task 6: Methodologies for Legal Evaluation](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Dhanachandra Ningthoujam, Pinal Patel, Rajkamal Kareddula, and Ramanand Vangipuram. 2023. [ResearchTeam\\_HCN at SemEval-2023 Task 6: A knowledge enhanced transformers based legal NLP system](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1245–1253, Toronto, Canada. Association for Computational Linguistics.
- Yuri Noviello, Enrico Pallotta, Flavio Pinzarrone, and Giuseppe Tanzi. 2023. [TeamUnibo at SemEval-2023 Task 6: A transformer based approach to Rhetorical Roles prediction and NER in Legal Texts](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 275–284, Toronto, Canada. Association for Computational Linguistics.
- Shounak Paul, Arpan Mandal, Pawan Goyal, and Saptarshi Ghosh. 2022. [Pre-training transformers on indian legal text](#). *arXiv preprint arXiv:2209.06049*.
- Nishchal Prasad, Mohand Boughanem, and Taoufiq Dkaki. 2023. [IRIT\\_IRIS\\_C at SemEval-2023 Task 6: A Multi-level Encoder-based Architecture for Judgment Prediction of Legal Cases and their Explanation](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 686–692, Toronto, Canada. Association for Computational Linguistics.
- Juliano Rabelo, Randy Goebel, Mi-Young Kim, Yoshinobu Kano, Masaharu Yoshioka, and Ken Satoh. 2022. [Overview and Discussion of the Competition on Legal Information Extraction/Entailment \(COLIEE\) 2021](#). *The Review of Socionetwork Strategies*, 16(1):111–133.
- Juliano Rabelo, Mi-Young Kim, Randy Goebel, Masaharu Yoshioka, Yoshinobu Kano, and Ken Satoh. 2020. [Coliee 2020: methods for legal document retrieval and entailment](#). In *JSAI International Symposium on Artificial Intelligence*, pages 196–210. Springer.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). *arXiv preprint arXiv:1908.10084*.
- Nicolay Rusnachenko, Thanet Markchom, and Huizhi Liang. 2023. [nclu\\_team at SemEval-2023 Task 6: Attention-based Approaches for Large Court Judgment Prediction with Explanation](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 270–274, Toronto, Canada. Association for Computational Linguistics.
- M Saravanan, Balaraman Ravindran, and S Raman. 2008. [Automatic identification of rhetorical roles using conditional random fields for legal document summarization](#). In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-1*.
- Isabel Segura-Bedmar, Paloma Martínez, and María Herrero-Zazo. 2013. [SemEval-2013 task 9 : Extraction of drug-drug interactions from biomedical texts \(DDIExtraction 2013\)](#). In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 341–350, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Patchipulusu Sindhu, Diya Gupta, Sanjeevi Meghana, and Anand Kumar M. 2023. [NITK\\_LEGAL at SemEval-2023 Task 6: A Hierarchical based system for identification of Rhetorical Roles in legal judgements](#). In *Proceedings of the 17th International*

- Workshop on Semantic Evaluation*, pages 1154–1160, Toronto, Canada. Association for Computational Linguistics.
- Kushagri Tandon and Niladri Chatterjee. 2023. [LRL\\_NC at SemEval-2023 Task 6: Sequential Sentence Classification for Legal Documents Using Topic Modeling Features](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 143–149, Toronto, Canada. Association for Computational Linguistics.
- Fareen Tasneem, Tashin Hossain, Jannatun Naim, and Abu Nowshed Chy. 2023. [CSECU-DSG at SemEval-2023 Task 6: Segmenting Legal Documents into Rhetorical Roles via Fine-tuned Transformer Architecture](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Quang-Minh Tran and Xuan-Dung Doan. 2023. [VTCC-NER at SemEval-2023 Task 6: An Ensemble Pre-trained Language Models for Named Entity Recognition](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 415–419, Toronto, Canada. Association for Computational Linguistics.
- Vu Tran, Minh Le Nguyen, and Ken Satoh. 2019. Building legal case retrieval systems with lexical matching and summarization using a pre-trained phrase scoring model. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, ICAIL '19*, page 275–282, New York, NY, USA. Association for Computing Machinery.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *International Conference on Learning Representations*.
- Serena Villata, Michal Araszkiwicz, Kevin Ashley, Trevor Bench-Capon, L Karl Branting, Jack G Conrad, and Adam Wyner. 2022. Thirty years of artificial intelligence and law: the third decade. *Artificial Intelligence and Law*, pages 1–31.
- Hannes Westermann, Jaromir Savelka, Vern R Walker, Kevin D Ashley, and Karim Benyekhlef. 2022. Data-centric machine learning in the legal domain. *arXiv preprint arXiv:2201.06653*.
- Nuo Xu, Pinghui Wang, Long Chen, Li Pan, Xiaoyan Wang, and Junzhou Zhao. 2020. Distinguish confusing law articles for legal judgment prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3086–3095.
- 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.
- Wenmian Yang, Weijia Jia, Xiaojie Zhou, and Yutao Luo. 2019a. Legal judgment prediction via multi-perspective bi-feedback network. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 4085–4091. International Joint Conferences on Artificial Intelligence Organization.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019b. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Junzhe Zhao, Yingxi Wang, Nicolay Rusnachenko, and Huizhi Liang. 2023. [Legal\\_try at SemEval-2023 Task 6: Voting Heterogeneous Models for Entities identification in Legal Documents](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 1282–1286, Toronto, Canada. Association for Computational Linguistics.
- Haoxi Zhong, Yuzhong Wang, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. Iteratively questioning and answering for interpretable legal judgment prediction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 1250–1257.