

# Leveraging Task Dependency and Contrastive Learning for Case Outcome Classification on European Court of Human Rights Cases

Santosh T.Y.S.S, Marcel Perez San Blas, Phillip Kemper, Matthias Grabmair

School of Computation, Information, and Technology;

Technical University of Munich, Germany

{santosh.tokala, marcel.perez, phillip.kemper, matthias.grabmair}@tum.de

## Abstract

We report on an experiment in case outcome classification on European Court of Human Rights cases where our model first learns to identify the convention articles allegedly violated by the state from case facts descriptions, and subsequently uses that information to classify whether the court finds a violation of those articles. We assess the dependency between these two tasks at the feature and outcome level. Furthermore, we leverage a hierarchical contrastive loss to pull together article-specific representations of cases at the higher level, leading to distinctive article clusters. The cases in each article cluster are further pulled closer based on their outcome, leading to sub-clusters of cases with similar outcomes. Our experiment results demonstrate that, given a static pre-trained encoder, our models produce a small but consistent improvement in classification performance over single-task and joint models without contrastive loss.

## 1 Introduction

The NLP task of classifying case outcome information from a textual statement of case facts is generally referred to as Legal Judgement Prediction (LJP)(e.g., Aletras et al. 2016; Chalkidis et al. 2019). It has been studied using corpora from different jurisdictions, such as the European Court of Human Rights (ECtHR) (Chalkidis et al., 2019, 2021, 2022b; Aletras et al., 2016; Medvedeva et al., 2020; Santosh et al., 2022, 2023), Chinese Criminal Courts (Luo et al., 2017; Zhong et al., 2018; Yue et al., 2021; Zhong et al., 2020; Yang et al., 2019), US Supreme Court (Katz et al., 2017; Kaufman et al., 2019), Indian Supreme Court (Malik et al., 2021; Shaikh et al., 2020), the French court of Cassation (Şulea et al., 2017), Brazilian courts (Bertalan and Ruiz, 2020), the Federal Supreme Court of Switzerland (Niklaus et al., 2021), the Turkish Constitutional court (Sert et al., 2021; Mumcuoğlu

et al., 2021), UK courts (Strickson and De La Iglesia, 2020), German courts (Waltl et al., 2017), the Philippine Supreme court (Virtucio et al., 2018), and the Thailand Supreme Court (Kowsrihawat et al., 2018).

In this work, we experiment with classifying case outcomes in the ECtHR A and B benchmark tasks introduced by LexGLUE (Chalkidis et al., 2022b). Task B is to identify the set of articles of the European Convention of Human Rights (ECHR) that the claimant alleges to have been violated, while Task A’s goal is to classify which of the convention’s articles has been deemed violated by the court. The input for both tasks is the case’s fact description that has been extracted from the published judgement document. It should be noted that, despite these tasks being typically referred to as instances of ‘legal judgement prediction’, the fact statements are typically not finalized until the decision outcome is known, making the task effectively one of retrospective classification rather than prediction (Medvedeva et al., 2021). While this does lead to distracting and confounding phenomena (see our prior work in Santosh et al. 2022), the dataset remains a useful resource for the development of NLP models that analyze these fact statements for text patterns that correspond to specific convention articles as drafted by the court. Consequently, we speak of *case outcome classification (COC)*.

Positive instances for Task A are cases in which an article was deemed violated by the court. Negative instances, however, usually encompass not only cases in which that particular article was alleged and considered not violated, but also the cases in which the particular article was not alleged in the first place. Given that the conditional probability of a positive task A label (violation) given its task B label (allegation) can be very high in the LexGLUE dataset, we posit that models can fall into this pitfall of identifying dominant articles

with high conditional violation probability, missing information specific to violation classification. Thus, we experiment with multi-task architectures to decouple these reasoning steps involved in Task A. This is similar to work in Chinese criminal cases judgement prediction carried out in [Zhong et al. 2018](#); [Yang et al. 2019](#); [Xu et al. 2020](#); [Yue et al. 2021](#) that coordinated multiple outcome variables (law articles, criminal charges, and penalty terms).

In the ECtHR context, most of the previous works ([Chalkidis et al., 2019, 2021, 2022b](#); [Clavié and Alphonsus, 2021](#); [Santosh et al., 2022, 2023](#)) employed independent models for allegation identification and violation classification. Concurrent work by [Valvoda et al. 2023](#) has recently explored multi-task joint models that learn Task A and Task B simultaneously. Their *Claim-Outcome model* decomposes them into two independent classifiers with separate encoders for each allegation and violation classification given allegation information. Our work differs in three ways: (i) We learn the joint representation for both tasks through a shared encoder following our intuition that the shared representation is beneficial as both tasks require similar features at the lower level; (ii) We model dependency from allegation to violation classification at both the allegation outcome and feature levels. Conditioning violation classification on the allegation outcome leads to higher performance, but we find that passing along feature level information yields additional improvement. This way, the violation branch will focus on identifying specific information required for determining violation than falling prey to only allegation level features. (iii) We also model correlations among different articles that tend to be concurrently alleged.

Inspired by recent advances in contrastive learning to learn effective representations ([Khosla et al., 2020](#)), we further devise a two-level hierarchical contrastive loss for COC. First, we strive to maximize the latent space distance between different article representations, which assists the model in learning distinct article-specific views of case facts. Second, we apply contrastive learning within each article latent space to form distinct sub-clusters of similar outcome cases. Unlike our two-level hierarchical contrastive learning, a single-level one has been explored for LJP in the concurrent work by [Zhang et al. 2023](#).

Our experiments demonstrate that, given a static

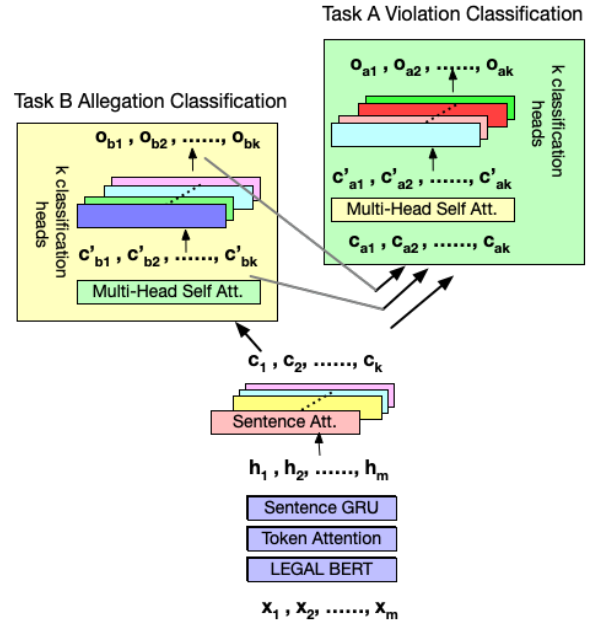


Figure 1: Our architecture capturing task dependency from allegation to violation classification branch.

pre-trained encoder, our models outperform single-task and joint models without contrastive loss by a small but consistent margin, with larger improvements for sparse classes in classification performance.

## 2 Method

Our model takes as input the case fact description as token encodings  $x = \{x_1, x_2, \dots, x_m\}$  where  $x_i = \{x_{i1}, x_{i2}, \dots, x_{in}\}$  and outputs the set of articles claimed to be alleged by the applicant (Task B) and set of articles deemed to be violated by the court (Task A) as multi-hot vectors.  $x_i$  and  $x_{ip}$  denote  $i^{\text{th}}$  sentence and  $p^{\text{th}}$  token of  $i^{\text{th}}$  sentence of the case facts respectively.  $m$  and  $n$  denote the number of sentences and tokens in the  $i^{\text{th}}$  fact sentence, respectively. Our model is built on top of a hierarchical attention network ([Yang et al., 2016](#)), which we chose to be able to process the very long fact description inputs. Our model contains three main components: (1) encoding layer, (2) Article-Specific Case representation and (3) Task Dependency learning. See Fig. 1 for an overview of our model.

**Encoding layer:** We encode each sentence in the case facts description with LegalBERT ([Chalkidis et al., 2020](#)) to obtain the token level representations  $\{z_{i1}, z_{i2}, \dots, z_{in}\}$ . Sentence level representations are then computed using a token-level atten-

tion mechanism as follows:

$$u_{it} = \tanh(W_w z_{it} + b_w) \ \& \ \alpha_{it} = \frac{\exp(u_{it} u_w)}{\sum_t \exp(u_{it} u_w)} \ \& \ f_i = \sum_{t=1}^n \alpha_{it} z_{it} \quad (1)$$

where  $W_w$ ,  $b_w$  and  $u_w$  are trainable parameters. The sentence level representations  $\{f_1, \dots, f_m\}$  are passed through a GRU encoder to obtain context-aware representations of the case facts  $h = \{h_1, h_2, \dots, h_m\}$ .

**Article-Specific Case Representation:** We disentangle the input case facts into multiple article-view-representations  $c = \{c_1, c_2, \dots, c_k\}$  where  $k$  is the total number of modeled convention articles, through aggregating context-aware sentence-level representations  $h$  using  $k$  sentence-level attention mechanisms similar to eq. 1 for every article individually. Through this article-specific attention mechanism, sentences relevant to a specific article are emphasized and intended to aid in fine-grained reasoning of outcome prediction for every article. This is helpful especially for sparser classes and mitigates the tendency of models to focus on skewed dominant articles. This is distinct from previous works (Chalkidis et al., 2019, 2021, 2022b; Santosh et al., 2022) which use a single vector representation for case facts.

**Task Dependency Learning:** To capture both the inter-article correlations for allegation, violation classification, and the inter-task dependency, we again use a two-step architecture. First, we apply a multi-head self-attention layer (Vaswani et al., 2017) to the obtained article-specific representations of case  $c$  to allow interactions among articles (e.g., some articles are typically alleged together while others usually occur in isolation) and obtain article-interaction-aware allegation feature representations  $c'_b = \{c'_{b1}, c'_{b2}, \dots, c'_{bk}\}$ . These  $c'_b$  are passed through  $k$  article-specific Task B classification layers to obtain the binary outcome  $o_b$  corresponding to each article for allegation classification, which are then concatenated into a multi-hot vector.

Then we utilize the task B allegation information to enhance task A violation classification to capture the task dependency. We concatenate the obtained article-interaction-aware allegation representations  $c'_b$  and allegation label probability logits  $o'_b$ <sup>1</sup> with article-specific case representations  $c$  to obtain the enhanced article-specific representation for Task A as  $c_{ai} = [c_i, c'_{bi}, o'_{bi}]$ . We again em-

<sup>1</sup>We use non-binarized response (i.e. probability logits) as it avoids the information loss that can occur in the probability space due to binarization.

ploy multi-head self-attention mechanism on these enhanced article-aware representations  $c_a$  for violation prediction to capture the correlations that exist between violation of different articles. Finally, the obtained representations are passed through  $k$  article-specific Task A classification layers to obtain the binary outcome  $o_a$  corresponding to each article’s violation output. To be able to evaluate the effectiveness of our interaction architecture in a clean way, we freeze the LegalBERT encoder weights in all our experiments.

## 2.1 Hierarchical Contrastive Loss

Contrastive learning has recently gained attention as a technique to obtain effective representations. In essence, it involves pulling together an ‘anchor point’ and its related samples while pushing it away from unrelated samples in the embedding space. Originally developed in self-supervised learning (Chen et al., 2020; Henaff, 2020), it has since been adopted in supervised settings (Khosla et al., 2020) where samples with the same/different labels are deemed related/unrelated with respect to an anchor.

In this work, we use a hierarchical contrastive loss (Liang et al., 2022) alongside the standard binary cross entropy loss on the task outcome probability. On the higher level, this is intended to form distinctive clusters of article-specific case representations. We hypothesize that this distinctiveness maximization constraint in turn helps the article-specific representation component to extract salient information with respect to each article more effectively. At the lower level, inside the latent space for each article, we further perform contrastive learning among cases based on their outcome for tasks A and B, respectively. This allows the positive outcome representations of cases under a specific article to stay closer and separate from the negative outcome cases, leading to formation of sub-clusters. We apply contrastive losses for Task A and Task B separately on the interaction-aware representations. It is calculated as the mean loss computed based on every article based representation as an anchor point. The loss for the interaction-aware representation of the  $j^{\text{th}}$  case specific to the  $i^{\text{th}}$  article for Task B  $c'_{bi_j}$  as an anchor point is calculated as follows:

$$l(c'_{bi_j}) = \log\left(\frac{\exp(c'_{bi_j} \cdot f / \tau_a)}{\sum_{g \in P(i,j)} \exp(c'_{bi_j} \cdot g / \tau_a)}\right) \times \alpha \sum_{f \in R(i,j)} \frac{\exp(c'_{bi_j} \cdot f / \tau_c)}{\sum_{g \in Q(i,j)} \exp(c'_{bi_j} \cdot g / \tau_c)} \quad (2)$$

where  $P(i, j) = \{c'_{bl_m} | l \neq i \ \& \ m \neq j\}$  (all rep-

representations except the anchor point),  $Q(i, j) = \{c'_{bl_m} | l = i \ \& \ m \neq j\}$  (all representations which share the same article as the anchor point),  $R(i, j) = \{c'_{bl_m} | l = i \ \& \ y_{i_j} = y_{l_m} \ \& \ m \neq j\}$ ,  $y_{i_j}$  the denotes binary outcome label of case  $j$  with respect to article  $i$  (all representations which share the same article and outcome of that article as the anchor point).  $\tau_a$  and  $\tau_c$  are scalar temperature parameters that control the penalties on negative samples. The first term in the above equation 2 denotes the contrastive loss among article representations (i.e., for a given anchor point, positive pairs are obtained by article-aware representations of cases in the batch which share the same article with the anchor point; negative pairs are the remaining article representations). The second term contrasts the cases based on their task outcome within specific article representations. Contrastive learning has shown to be more effective with larger batch sizes (Radford et al., 2021; Chen et al., 2020). To account for smaller batch size due to computational constraints incurred with our hierarchical setup, we use a memory bank (Wu et al., 2018) which progressively reuses the representations from previous batches in computing the contrastive loss.

### 3 Experiments & Discussion

#### 3.1 Dataset & Baseline Models

We experiment on the ECHR task A and B datasets of LexGLUE (Chalkidis et al., 2022b), which consist of 11k case fact descriptions chronologically split into training (2001–2016, 9k cases), validation (2016–2017, 1k cases), and test sets (2017–2019, 1k cases). The label set for both tasks include 10 prominent ECHR articles. Implementation details are given in Appendix A. We compare our models against the following baselines:

**Task A/B only:** We train single-task models only on task A and B labels, respectively. Their architecture is based on our model with contrastive loss, task dependency and article-specific representation components removed.

**Multi-task:** Using the same architecture, we also develop a model for both task A and B that is trained in classic multi-task fashion with separate classification heads without article-specific representations and task dependency components.

#### 3.2 Performance Evaluation

Following Chalkidis et al. 2022b, we report micro-F1 ( $\mu$ -F1) and macro-F1 (m-F1) scores across the

Table 1: Results on Task A and B. ‘feat.’, ‘lab.’, ‘Cont.’ indicate feature, label and contrastive, respectively

Model	Task B		Task A		
	$\mu$ -F1	m-F1	$\mu$ -F1	m-F1	hm-F1
Task B only	76.20	67.15	-	-	-
Task A only	-	-	68.42	56.26	54.14
Multi-task	78.17	69.16	69.29	58.05	55.57
Our Method	79.29	70.97	71.26	65.24	60.90
w/o feat. & lab.	78.93	71.45	70.07	59.14	57.09
w/o feat.	78.59	71.56	70.68	63.93	59.28
w/o label	79.09	71.38	70.32	64.12	61.70
only gold lab.	78.21	70.03	81.46	78.93	66.59
gold lab. + feat.	77.68	70.40	83.19	78.79	67.42
w/o outcome Cont.	78.42	69.48	69.86	60.84	57.62
w/o article Cont.	79.02	71.14	71.16	64.68	59.86
Task A Cont.	-	-	70.16	62.14	58.12
Task B Cont.	78.16	69.42	-	-	-

10 ECtHR articles contained in the dataset. We also report hard-macro-F1 (hm-F1) for Task A following Santosh et al. 2022, which is the mean F1-score computed for each article where cases with that article having been violated are considered as positive instances, and cases with that article being alleged but not found to have been violated as negative instances. This forms the most critical measure for violation classification as it conditions on allegation.

We compare the performance of our method with the baselines in the first four rows of Table 1. The multi-task method performs better than the individual tasks alone, validating the dependency between them. Our method performs better than the multi-task architecture, highlighting the effectiveness of our feature and label dependency components, as well as of the hierarchical contrastive loss. It scores higher by a small margin in Task B and Task A micro-F1 but achieves a larger benefit (>5%) in task A macro-F1 and hard macro-F1 metrics. This suggests that our model is of particular utility for sparser classes.

#### 3.3 Analysis

##### 3.3.1 Ablation on Task Dependency Learning

We create three ablation conditions by removing one interaction mechanism in each of them: (i) *w/o feature & label* (article-specific representation but with no concatenation with features or labels) (ii) *w/o feature* (article-specific representation concatenated with task B label as classified by the model) and (iii) *w/o label* (article-specific representation concatenated with task B features).

From the second section of Table 1, we observe that even the performance of *w/o feature and label* is better than *multi-task* model indicated by a small but consistent margin, indicating that article-

specific case representation is a competitive component. Both *w/o feature*, *w/o label* models perform better than *w/o feature & label* in Task A, demonstrating the benefit of passing on task B model information explicitly. Between them, *w/o label* performs better in m-F1 and hm-F1 scores of Task A than *w/o feature* showing that providing the fine-grained representation of features from the allegation identification model is more useful than the predicted allegation labels only.

Further, to evaluate the impact of allegation identification performance on downstream violation classification, we conduct a control experiment in which we provide the actual gold allegation set of articles as input to the violation component in place of the predicted ones. We create two variants: *only gold labels* and *gold labels + features* (gold labels along with task B features). From the third section of Table 1, we observe that performance on Task A increases substantially in micro-F1 and macro-F1, which is intuitive as the model gets access to perfect allegation information. Interestingly, when adding task allegation feature information, we notice an additional small increase in hard macro-F1, indicating that adding feature information can have beneficial effect even in the presence of gold task B labels. These results form an upper bound of the benefit of using task dependency information from allegation branch to violation branch in our architecture, and the size of the performance gap motivates future work on accurate allegation classification.

### 3.3.2 Ablation Hierarchical Contrastive Loss

We carry out an ablation experiment for each component of our hierarchical contrastive loss: (i) disable article-level contrastive learning to obtain distinctive article representations (*‘w/o article contrastive’*) and (ii) disable outcome-based contrastive learning to separate cases based on outcome within each article cluster (*‘w/o outcome contrastive’*). Intuitively, from the fourth section of Table 1, we observe that removing the outcome contrastive loss has the larger effect on performance, as it directly relates to the predicted label. The smaller but consistent drop in performance when removing article-level contrastive learning supports our hypothesis that maximizing the distinctiveness among article representations encourages the model to learn how to extract article-specific salient information from case facts.

Finally, to study the impact of our contrastive

loss component alone, we evaluate single-task models applying our hierarchical contrastive loss referred to as *‘Task A/B Contrastive’*. From the last section of Table 1, we observe that they outperform their baseline counterparts in both tasks, but still stay behind our model.

## 4 Conclusion

We improve ECtHR article violation classification from fact statements by leveraging feature and label information from an allegation classification model. We also leveraged hierarchical contrastive loss to contrast between different article representations and case representations based on outcome with respect to a specific article. Given a static pre-trained encoder, our models outperform a straightforward multi-task architecture by a small but consistent margin, with larger improvements for sparse classes. These results suggest that the tasks of allegation and violation classification on ECtHR fact statements interrelate in a way that may not be optimally captured using straightforward multi-task architectures, and motivate further research on dependency modeling between related legal classification tasks.

## Limitations

In this work, we have demonstrated improvements in violation classification for ECtHR cases by leveraging the dependency from allegation classification at the feature and label level. While our feature transfer and contrastive learning techniques are general, our experimental contributions are contextual to the court. The nature of the fact statements as being influenced by the eventual case outcome and not suitable for prospective prediction has already been addressed in section 1.

While similar tasks (i.e., allegations paired with findings) exist in many other jurisdictions, they will differ in, for example, legal nature, semantic difficulty, and sub-task dependency. In our case, we choose the allegation classification as the auxiliary task which is closely related and also forms a necessary sub-part of deriving the final outcome related to the main task (violation prediction) at hand. We leave an exploration of the relatedness of auxiliary tasks and their impact on LJP/COC for future work.

One major challenge is dealing with long input case facts description, which is currently handled with hierarchical model in this work. These hierar-

chical models do not allow tokens in one sentence to attend to tokens in other sentence which leads to sub-optimal interaction modelling. This modeling impact on performance is still underexplored except some preliminary empirical work in Dai et al. 2022; Chalkidis et al. 2022a. Additionally, we freeze the weights in the LegalBERT encoder, both to save computational resources and to evaluate the effectiveness of our dependency mechanism in a clean way.

## Ethics Statement

We employ pre-trained language models and do not train them from scratch, thus inheriting the biases they may have acquired from their training corpus. Our experiments have been carried out on a dataset of ECtHR decisions which is publicly available as the part of LexGLUE benchmark (Chalkidis et al., 2022b) and has been derived from the public court database HUDOC<sup>2</sup>. Though these decisions are not anonymized and contain the real names of the involved parties, we do not foresee any harm incurred by our experiments beyond making this information available. This collection of decision documents is of course historical data and using it to train model may result in classifiers that exhibit behavior that may be considered biased. For example, Chalkidis et al. 2022c explores disparities in classification performance with regard to an applicant’s gender, their age, and the identity of the respondent state. If COC models are deployed as part of a decision support systems, then they of course must be screened for performance/error differences in between groups that are to be treated equally.

The task of LJP/COC in itself raises serious ethical and legal concerns, both in general and specific to the European Court of Human Rights. However, we do not advocate for the practical adoption of LJP/COC systems by courts. Our prior work in Santosh et al. 2022 demonstrates that these systems rely on several shallow surface-level spurious signals that are statistically predictive but legally irrelevant. This highlights the risk of using predictive systems in high stakes domains such as law. In the same work, we argue that models leveraging the case outcome signal for analytical purposes must be developed mindfully and with the goal of aligning their inferences with legal expert reasoning. This further parallels the broader legal NLP

community increasingly addressing ethical aspects of developed systems in the context of technical research (e.g., Wang et al. 2021; Medvedeva et al. 2021, 2022; Tsarapatsanis and Aletras 2021; Leins et al. 2020).

In this work, we use COC as a technical benchmarking task that allows the development and study of neural NLP models on legal text. We focus on how to leverage dependencies on two successive tasks (allegation identification and violation classification) based on case facts, as well as on learning effective representations of these facts using contrastive learning. Our results are hence to be understood as technical contributions in pursuit of the overarching goal of developing models capable of deriving insight from data that can be used legally, ethically, and mindfully by experts in solving problems arising in legal research and practice.

All experiments were carried out using Google Colab. We did not track computation hours.

## References

- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel PreoŃiuc-Pietro, and Vasileios Lampos. 2016. Predicting judicial decisions of the european court of human rights: A natural language processing perspective. *PeerJ Computer Science*, 2:e93.
- Vithor Gomes Ferreira Bertalan and Evandro Eduardo Seron Ruiz. 2020. Predicting judicial outcomes in the brazilian legal system using textual features. In *DHandNLP@ PROPOR*, pages 22–32.
- 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.
- Ilias Chalkidis, Xiang Dai, Manos Fergadiotis, Prodromos Malakasiotis, and Desmond Elliott. 2022a. An exploration of hierarchical attention transformers for efficient long document classification. *arXiv preprint arXiv:2210.05529*.
- 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.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. Paragraph-level rationale extraction through regularization: A case study on european court of human rights cases. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational*

<sup>2</sup><https://hudoc.echr.coe.int>

- Linguistics: Human Language Technologies*, pages 226–241.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022b. Lexglue: A benchmark dataset for legal language understanding in english. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330.
- Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. 2022c. Fairlex: A multilingual benchmark for evaluating fairness in legal text processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4389–4406.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Benjamin Clavié and Marc Alphonse. 2021. The unreasonable effectiveness of the baseline: Discussing svms in legal text classification. In *Legal Knowledge and Information Systems*, pages 58–61. IOS Press.
- Xiang Dai, Ilias Chalkidis, Sune Darkner, and Desmond Elliott. 2022. [Revisiting transformer-based models for long document classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7212–7230, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Olivier Henaff. 2020. Data-efficient image recognition with contrastive predictive coding. In *International conference on machine learning*, pages 4182–4192. PMLR.
- Daniel Martin Katz, Michael J Bommarito, and Josh Blackman. 2017. A general approach for predicting the behavior of the supreme court of the united states. *PLoS one*, 12(4):e0174698.
- Aaron Russell Kaufman, Peter Kraft, and Maya Sen. 2019. Improving supreme court forecasting using boosted decision trees. *Political Analysis*, 27(3):381–387.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: a method for stochastic optimization 3rd int. In *International Conference on Learning Representations*.
- Kankawin Kowsrihawat, Peerapon Vateekul, and Prachya Boonkwan. 2018. Predicting judicial decisions of criminal cases from thai supreme court using bi-directional gru with attention mechanism. In *2018 5th Asian Conference on Defense Technology (ACDT)*, pages 50–55. IEEE.
- Kobi Leins, Jey Han Lau, and Timothy Baldwin. 2020. Give me convenience and give her death: Who should decide what uses of nlp are appropriate, and on what basis? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2908–2913.
- Bin Liang, Zixiao Chen, Lin Gui, Yulan He, Min Yang, and Ruifeng Xu. 2022. Zero-shot stance detection via contrastive learning. In *Proceedings of the ACM Web Conference 2022*, pages 2738–2747.
- Bingfeng Luo, Yansong Feng, Jianbo Xu, Xiang Zhang, and Dongyan Zhao. 2017. Learning to predict charges for criminal cases with legal basis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2727–2736.
- 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.
- Masha Medvedeva, Ahmet Üstün, Xiao Xu, Michel Vols, and Martijn Wieling. 2021. Automatic judgement forecasting for pending applications of the european court of human rights. In *ASAIL/LegalAIIA@ICAIL*.
- Masha Medvedeva, Michel Vols, and Martijn Wieling. 2020. Using machine learning to predict decisions of the european court of human rights. *Artificial Intelligence and Law*, 28(2):237–266.
- Masha Medvedeva, Martijn Wieling, and Michel Vols. 2022. Rethinking the field of automatic prediction of court decisions. *Artificial Intelligence and Law*, pages 1–18.
- Emre Mumcuoğlu, Ceyhan E Öztürk, Haldun M Ozaktas, and Aykut Koç. 2021. Natural language processing in law: Prediction of outcomes in the higher courts of turkey. *Information Processing & Management*, 58(5):102684.
- Joel Niklaus, Ilias Chalkidis, and Matthias Stürmer. 2021. Swiss-judgment-prediction: A multilingual legal judgment prediction benchmark. In *Proceedings of the Natural Legal Language Processing Workshop 2021*, pages 19–35.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.

- T. Y. S. S Santosh, Oana Ichim, and Matthias Grabmair. 2023. Zero shot transfer of article-aware legal outcome classification for european court of human rights cases. *arXiv preprint arXiv:2302.00609*.
- T.y.s.s Santosh, Shanshan Xu, Oana Ichim, and Matthias Grabmair. 2022. [Deconfounding legal judgment prediction for European court of human rights cases towards better alignment with experts](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1120–1138, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mehmet Fatih Sert, Engin Yıldırım, and İrfan Haşlak. 2021. Using artificial intelligence to predict decisions of the turkish constitutional court. *Social Science Computer Review*, page 08944393211010398.
- Rafe Athar Shaikh, Tirath Prasad Sahu, and Veena Anand. 2020. Predicting outcomes of legal cases based on legal factors using classifiers. *Procedia Computer Science*, 167:2393–2402.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Benjamin Strickson and Beatriz De La Iglesia. 2020. Legal judgement prediction for uk courts. In *Proceedings of the 2020 the 3rd international conference on information science and system*, pages 204–209.
- Octavia-Maria Şulea, Marcos Zampieri, Mihaela Vela, and Josef van Genabith. 2017. Predicting the law area and decisions of french supreme court cases. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 716–722.
- Dimitrios Tsarapatsanis and Nikolaos Aletras. 2021. On the ethical limits of natural language processing on legal text. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3590–3599.
- Josef Valvoda, Ryan Cotterell, and Simone Teufel. 2023. On the role of negative precedent in legal outcome prediction. *Transactions of the Association for Computational Linguistics*, 11:34–48.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Michael Benedict L Virtucio, Jeffrey A Aborot, John Kevin C Abonita, Roxanne S Avinante, Rother Jay B Copino, Michelle P Neverida, Vanesa O Osiana, Elmer C Peramo, Joanna G Syjuco, and Glenn Brian A Tan. 2018. Predicting decisions of the philippine supreme court using natural language processing and machine learning. In *2018 IEEE 42nd annual computer software and applications conference (COMPSAC)*, volume 2, pages 130–135. IEEE.
- Bernhard Waltl, Georg Bonczek, Elena Scepankova, Jörg Landthaler, and Florian Matthes. 2017. Predicting the outcome of appeal decisions in germany’s tax law. In *International conference on electronic participation*, pages 89–99. Springer.
- Yuzhong Wang, Chaojun Xiao, Shirong Ma, Haoxi Zhong, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2021. Equality before the law: Legal judgment consistency analysis for fairness. *arXiv e-prints*, pages arXiv–2103.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.
- 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*, pages 3086–3095.
- Wenmian Yang, Weijia Jia, Xiaojie Zhou, and Yutao Luo. 2019. Legal judgment prediction via multi-perspective bi-feedback network. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 4085–4091.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489.
- Linan Yue, Qi Liu, Binbin Jin, Han Wu, Kai Zhang, Yanqing An, Mingyue Cheng, Biao Yin, and Dayong Wu. 2021. Neurjudge: a circumstance-aware neural framework for legal judgment prediction. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 973–982.
- Han Zhang, Zhicheng Dou, Yutao Zhu, and Ji-Rong Wen. 2023. Contrastive learning for legal judgment prediction. *ACM Transactions on Information Systems*.
- Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. Legal judgment prediction via topological learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3540–3549.



Haoxi Zhong, Yuzhong Wang, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. Iteratively questioning and answering for interpretable legal judgment prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1250–1257.

## A Implementation Details

Our models compute word embeddings of size 768. Our word level attention context vector size is 300. The sentence level GRU encoder dimension is 200, thus giving a bidirectional embedding of size 400, and a sentence level attention vector dimension of 200. The final dense classifier for all tasks has 100 hidden units. We use a mini batches size of 32 and the model is optimized end-to-end using Adam (Kingma and Ba, 2015). The dropout rate (Srivastava et al., 2014) in all layers is 0.1. We determine the best learning rate using grid search on the development set and use early stopping based on the development set m-F1 score. We finetuned  $\tau_a$ ,  $\tau_c$  with an additional constraint of  $\tau_a < \tau_c$  among the values of  $\{0.07, 0.1, 0.14, 0.2, 0.25, 0.3\}$  so that it aids in pulling together the representations belonging to the same article in latent space (leading to distinct article clusters) and also in further slightly pulling together the representations of cases belonging to the same outcome in each separated article-specific embedding latent space compared to the other outcome cases in that same article. We set  $\alpha$  to be 0.5. We use a memory bank of size 32 per article and outcome, and store only the most recent examples per article and its corresponding outcome.