

Towards Explainability in Legal Outcome Prediction Models

Josef Valvoda

University of Cambridge
jv406@cam.ac.uk

Ryan Cotterell

ETH Zürich
ryan.cotterell@inf.ethz.ch

Abstract

Current legal outcome prediction models—a staple of legal NLP—do not explain their reasoning. However, to employ these models in the real world, human legal actors need to be able to understand the model’s decisions. In the case of common law, legal practitioners reason towards the outcome of a case by referring to past case law, known as precedent. We contend that precedent is, therefore, a natural way of facilitating explainability for legal NLP models. In this paper, we contribute a novel method for identifying the precedent employed by legal outcome prediction models. Furthermore, by developing a taxonomy of legal precedent, we are able to compare human judges and neural models with respect to the different types of precedent they rely on. We find that while the models learn to predict outcomes reasonably well, their use of precedent is unlike that of human judges.

 https://github.com/valvoda/under_the_influence

1 Introduction

Legal outcome prediction models have generated much interest in the past years (Aletras et al., 2016; Medvedeva et al., 2020; T.y.s.s et al., 2023; Valvoda et al., 2023a). Given a text transcript of the facts of a legal case, a legal outcome prediction model predicts whether a particular law has been breached. Researchers have evaluated these models in several jurisdictions, including the UK, USA, Switzerland, France, China, Japan, India and the European Court of Human Rights (ECTHR; Şulea et al., 2017; Chalkidis et al., 2019; Strickson and De La Iglesia, 2020; Malik et al., 2021; Niklaus et al., 2021; Chalkidis et al., 2022; Feng et al., 2022).

While legal outcome prediction models are interesting in the context of academic research (Valvoda et al., 2021), any real-world deployment of Legal AI necessitates a certain level of understanding of how the system works (Valvoda et al., 2023b). After all, it is the *reasoning* towards the outcome of the case, rather than the outcome itself, that is of tantamount importance in legal outcome prediction.

There are two reasons for this. (1) To be able to use a legal outcome prediction model system in practice implicitly assumes explainability: lawyers are bound to provide advice that the client can depend on. This obligation is so extreme that a negligent lawyer themselves can be sued for misconduct if they give faulty advice. Moreover, making legal decisions without understanding the reasoning behind them is unethical in many legal systems. For example, for a model to be used in practice it is a legal requirement in the EU under GDPR to know what the decision is based on. Thus, the ability to double-check any information a hypothetical legal AI assistant provides is paramount when building technology for legal professionals. (2) Explainability can create new use cases for legal outcome prediction models. Knowing how to reason towards a desired legal outcome is valuable for drafting legal arguments and, thus, a legal outcome prediction system could be deployed as an assistive technology in areas well beyond simply predicting outcomes (Sarker et al., 2019).

State-of-the-art legal outcome prediction models are nearly homogeneous at a technical level—they are all parameterized as neural network classifiers on top of pre-trained language models (LM). While their performance is remarkable at times, the reasoning embedded in the model is largely inscrutable. Despite the widespread recognition of the problem that unexplainable legal NLP systems pose (Hacker et al., 2020) and emerging work in the domain of case law summarization (Norkute et al., 2021), much remains to be done in order to build effective explainable legal NLP systems (Branting et al., 2019; Chalkidis et al., 2021). In particular, the existing work on explainable legal AI focuses on providing explanations by identifying relevant tokens (Xu et al., 2023), sentences (Malik et al., 2021) or paragraphs (Chalkidis et al., 2021), from the facts of the case. However, lawyers and judges typically reason at the level of a case.

In this paper, we address the lack of explainability of legal outcome prediction models in a manner inspired by legal reasoning. The law in the Anglo-American tradition is endowed with

a natural rationale—precedent—which produces a legal argument. Precedent is how judges and lawyers reason with respect to previous cases. Because precedent is binding, by citing previously decided cases, legal actors substantiate their arguments. Unlike existing explainability methods, our method produces an explanation in terms of legal precedent. This allows legal practitioners to easily inspect the legal reasoning the model employs. Furthermore, we develop a taxonomy of legal precedent, which allows us to categorize each training case in terms of its precedential operation. We can therefore pinpoint which kinds of precedential reasoning the models rely on.

We validate our proposed method by training several legal outcome prediction models and measuring the correlation between the precedent they use and that used by a human judge. Overall, our experiments reveal that the precedent used by a human judge has a weak positive correlation with the precedent our models rely on (the highest Spearman’s ρ we observe is 0.18). We also discover that the model predictions positively correlate with one type of precedent in particular—the one where the outcome of the precedent case is the same as the outcome of the case at hand. However, for most of the categories of precedent, we find the models have a negligible or even negative correlation with the precedent applied by a human judge. Worse yet, model F_1 performance on predicting the outcome of a case is a poor proxy for its alignment with precedential reasoning; see Tab. 2. In other words, higher F_1 does not equal better alignment. Therefore, our results cast doubt on the real-world utility of the current outcome prediction models.

2 Precedent-based Interpretability

We now give an argument for why precedent-based interpretability is the right direction for improving the usability of legal outcome prediction models.

Why are existing methods insufficient? Many interpretability methods, such as linguistic probing, have been developed to investigate the linguistic capabilities of neural models (Hewitt and Manning, 2019; Belinkov, 2022; Pimentel et al., 2022). These methods often focus solely on the immediate input to the model. This is intentional because the goal of linguistic probing is to analyze the model in terms of how individual tokens affect the output. For example, if you want to know if and how a neural network’s representations relate to

a dependency parse of the input sentence, investigating the network’s use of tokens in a sentence is a wise choice. In legal NLP, however, despite the recent rejuvenation of efforts in the field (Xu et al., 2023; Habernal et al., 2023), there is no generally accepted legal theory about how individual words or sentences affect the outcome of a case. We further elaborate on the related work in §9.

Why precedent? There is already a generally accepted method of legal reasoning—reasoning through legal precedent. However, precedent does not operate on the granularity of individual words or sentences. Instead, precedent relies on the relationship between the case at hand and previously decided cases. Importantly, precedent is a way of developing law. A decision rendered in a past case binds future decisions in a common law system. This binding precedent, which stems from the doctrine of *stare decisis*, literally *stand by the decision*, ensures the reliability of law by forcing judges to make consistent decisions (Duxbury, 2008; Lamond, 2016; Black, 2019).¹ While precedent plays an important role in deciding the outcome of a case in both common and civil law systems, it occupies the core of legal reasoning in common law jurisdictions, e.g., UK, USA, and India, and it is the focus of this paper.

Claims and Outcomes. Because precedent is binding law, it is used in legal argumentation. When a new legal case arises, there are two competing points of view presented to a judge: That of the plaintiff and that of the defendant. These are typically represented by lawyers working on behalf of either side. The first step in the process is establishing the **claims**, i.e., the laws that the plaintiff claims as violated. The claims are a starting point of the legal argument. The job of the lawyer on either side is to substantiate their claims. To do this, the lawyers rely on previous case law, i.e., the precedent. They will draw analogies to cases where their side has won, and distinguish their circumstances from the instances where their side has lost. The selection of past cases is therefore a large part

¹There is also a non-binding precedent. It is typically contained in *obiter dicta* or the *other things said* in the court decisions. Its purpose is to hint at what the court might decide in the future given a legal point that falls outside the scope of the case under consideration. Another type of non-binding precedent is in cases that have been overruled by a subsequent judgment. In such instances, the case is considered outdated and the law contained within is obsolete. While this kind of precedent is important, we do not study them in this paper.

of the reasoning process for the lawyer. Lawyers present their arguments, substantiated with citations to precedent cases, to the judge, who again uses precedent in their argumentation. The judge will cite past cases, much like lawyers do, as part of their argument towards the **outcome** of a case. As a judge arrives at the decision, the case at hand will set a binding precedent and the process repeats.

Conclusion. Our position is that the best way to build explainable legal outcome prediction models is through the direct use of precedent. This way, both lawyers and judges can directly interpret and analyze the rationale provided by the model. Moreover, both lawyers and judges are already trained in reasoning about new cases with respect to old cases, so no additional training would be required.

3 Under the Influence

We now turn to the explainability method developed in this paper, which will allow us to identify the exact role precedent plays in the decisions made by neural legal outcome prediction models.

3.1 Preliminaries

Legal outcome prediction is a multilabel classification task. We start with the necessary notation and a generic description of a probabilistic legal outcome prediction model. Let $\mathcal{O} = \{+, -, \emptyset\}$ be a set of **legal outcomes**, whose elements correspond to a positive, negative and null outcome of a case, respectively. The null outcome occurs when a law has not been claimed by either party to the case. Let O be an outcome-valued random variable with instances $o \in \mathcal{O}$. Let \mathcal{O}^K be the K -times Cartesian product of \mathcal{O} where K is the number of legal Articles we consider. The elements of \mathcal{O}^K are denoted as $\mathbf{o} \in \mathcal{O}^K$. Let \mathbf{O} is a random variable ranging over \mathcal{O}^K . We write O_k for the random variable, ranging over \mathcal{O} , that corresponds to just the k^{th} article. For instance, if a judge finds that Article 3 has been breached, we write $O_k = +$. Now, let Σ be an alphabet.² The random variable \mathbf{F} ranges over textual descriptions of **facts**, i.e., Σ^* for a vocabulary Σ . Elements of Σ^* are denoted as \mathbf{f} . For a case concerned with O_5 we might expect the following \mathbf{f} .

Example 1 (Article 5 - Right to a fair trial). *Neither the applicants nor any of the other migrants in their group were provided with any information*

²In principle, the alphabet Σ can be constructed over either characters, words or sub-word tokens, the last of which is what we do in our experiments.

regarding the reason for or the length of their detention. They did not have a chance to appoint a lawyer and were not provided with an interpreter at any point during their detention.

We focus on an **article-factored** probabilistic legal outcome predictor, which can be described by the following probability model

$$p_{\theta}(\mathbf{O} = \mathbf{o} \mid \mathbf{F} = \mathbf{f}) = \prod_{k=1}^K p_{\theta}(O_k = o_k \mid \mathbf{F} = \mathbf{f}) \quad (1)$$

parameterized by $\theta \in \Theta$. We further assume that Θ is a compact subset of \mathbb{R}^d . The training dataset $\mathcal{D}_{\text{trn}} = \{\mathbf{z}^{(n)}\}_{n=1}^N$ is a set of N outcome-fact pairs where $\mathbf{z}^{(n)} = (\mathbf{f}^{(n)}, \mathbf{o}^{(n)})$. The test dataset $\mathcal{D}_{\text{tst}} = \{\mathbf{z}^{(m)}\}_{m=1}^M$ is a set of M outcome-fact pairs where $\mathbf{z}^{(m)} = (\mathbf{f}^{(m)}, \mathbf{o}^{(m)})$. Each training and test pair encodes a relationship between facts and law. The instance-level negative log-likelihood, i.e., the negative log-likelihood of a *single* training point \mathbf{z} , is denoted as $\ell_{\mathbf{z}}(\theta) = -\log p_{\theta}(\mathbf{z})$. Then, the average cross-entropy for the training set is given by

$$L(\theta) = \frac{1}{N} \sum_{n=1}^N \ell_{\mathbf{z}^{(n)}}(\theta) \quad (2)$$

A good value for the parameters θ may be found by solving the following optimization problem

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} L(\theta) + \alpha \mathcal{R}(\theta) \quad (3)$$

where $\mathcal{R}(\cdot)$ is a regularization function with strength $\alpha \geq 0$. Because Θ is assumed to be compact and $L(\theta) + \alpha \mathcal{R}(\theta)$ is a continuous function of θ the argmin given in Eq. (3) is well-defined.

3.2 Influence Functions

In order to evaluate if and to what extent a model relies on any of the above categories of precedent, we need a way to identify how precedent cases affect the model's decision. In a legal outcome prediction model, the only cases that can be considered precedent are the training cases. A simple technique to measure the effect of a training case \mathbf{z} on a test case \mathbf{z}' would be to remove \mathbf{z} from the training set \mathcal{D}_{trn} , retrain the model, and take \mathbf{z} 's importance to be the change in the loss $\ell_{\mathbf{z}'}(\theta)$ for a test point \mathbf{z}' . However, with over $N = 8000$ training cases, we would have to retrain the model N times to be able to estimate the importance of each case in the training set for the test cases \mathbf{z}' . This would be prohibitively slow.

As a tractable alternative to the above, Koh and Liang (2017) introduce a simple procedure that allows us to efficiently estimate the effect of a single z on the model parameters *without* retraining. We now describe their method, based on influence functions, in detail. Consider the following optimization problem, a modified version of Eq. (3)

$$\theta_z^*(\varepsilon) = \operatorname{argmin}_{\theta \in \Theta} \underbrace{L(\theta) + \alpha \mathcal{R}(\theta)}_{\stackrel{\text{def}}{=} L_{\mathcal{R}}(\theta)} + \varepsilon \ell_z(\theta) \quad (4)$$

It is not hard to see that Eq. (4) simply modifies the weight associated with z in the summation found in Eq. (3). Relying on Koh and Liang’s (2017) method, we can approximate the effect of a very small ε on parameters θ by taking a derivative of $\theta_z^*(\varepsilon)$ with respect to ε at evaluating it at $\varepsilon = 0$:

$$\left. \frac{d\theta_z^*}{d\varepsilon} \right|_{\varepsilon=0} = - (\nabla_{\theta}^2 L_{\mathcal{R}}(\theta^*))^{-1} \nabla_{\theta} \ell_z(\theta^*) \quad (5)$$

See App. F for a derivation of Eq. (5).

Influence score. We are not directly interested in the new parameters themselves but rather in the effect of the change in parameters on the loss for a test point z' . We can calculate this influence score using the chain rule:

$$\begin{aligned} \iota(z, z') &\stackrel{\text{def}}{=} \left. \frac{d\ell_{z'}(\theta_z^*)}{d\varepsilon} \right|_{\varepsilon=0} \\ &= \nabla_{\theta} \ell_{z'}(\theta^*)^{\top} \left. \frac{d\theta_z^*}{d\varepsilon} \right|_{\varepsilon=0} \quad (\text{chain rule}) \quad (6) \\ &\stackrel{(a)}{=} -\nabla_{\theta} \ell_{z'}(\theta^*)^{\top} (\nabla_{\theta}^2 L_{\mathcal{R}}(\theta^*))^{-1} \nabla_{\theta} \ell_z(\theta^*) \end{aligned}$$

where (a) results by plugging in Eq. (5). Thus, we can leverage Koh and Liang’s method to measure the extent to which a precedent case bears on a new legal decision.

4 Taxonomy of Precedent

Precedent is more than just a binary category. Legal actors use it in a number of different ways to craft their arguments. To better evaluate how neural models of outcome use precedent, we develop a novel taxonomy of types of precedent.

4.1 Applied vs. Distinguished

The simplest way of categorizing the precedential relationship between two cases is to differentiate between application and distinction.³

³We focus on the major types of precedent. There are other forms of precedent we do not consider here. Specifically, modification of precedent and overruling. In future work, they would be a natural candidate for further exploration.

Applied Precedent. We call a precedent case **applied** when a judge applies the precedent from a previous case. In such instances, the judge has relied on similar reasoning to that of the precedent case. As a result, the outcome of the case at hand is the same as the outcome of the precedent case. In practice, these are the precedent cases that a judge has cited as part of their argument, which have the same outcome as the case at hand.

Distinguished Precedent. We call a precedent case **distinguished** when a judge distinguishes their decision from a past case. When a precedent is distinguished, the judge has used the differences between the cases to inform their decision. In practice, these are the cases that a judge cites but have a different outcome than the case at hand.

4.2 Positive vs. Negative Precedent

Valvoda et al. (2023a) introduce a new perspective on how to model legal outcome prediction. By considering both the claims and the outcomes of cases, they define two types of legal outcomes: positive and negative.⁴ Following their logic, we can further sub-divide the applied or distinguished precedent into the two categories below.

Positive Precedent. An outcome is positive when a law has been claimed as breached and a judge finds in favor of the claimant. When a new case cites a positive outcome case, the judge invokes a positive precedent. In the strictest sense, this means that if a new case arises from exactly the same facts as the precedent case before it, the judge will have to conclude the same law has been equally breached in both cases. In practice, no two cases are exactly the same. Therefore the positive precedent affects all subsequent similar cases by increasing the chances of successfully claiming a breach of the law that has been violated in the precedent case.

Negative Precedent. An outcome is negative when a law has been claimed as breached, but the judge finds it has not in fact been broken. When a new case cites a negative outcome case, the judge invokes a negative precedent. Much like with positive precedent, the court is bound by the negative precedent. Two cases with the same facts must reach the same outcome—positive or

⁴This is not a new concept to lawyers or legal AI theorists. Lawlor (1963) already described these categories of precedent as *pro-precedent* and *con-precedent*.

	Applied		Distinguished	
	Cited	Claimed	Cited	Claimed
Pos.	2,012	643,377	114	54,099
Neg.	296	129,239	79	30,203

Table 1: Total number of cited cases per category of precedent under the per-case view of the precedent.

negative. Therefore, the scope of the law contracts through negative precedent.

4.3 Scope

Additionally, there are two axes along which we can define the scope of the precedent we will consider in §7.

Claimed vs. Cited. There are two ways to uncover the relevant precedent for a case. The first is to consider all the Articles that the lawyer has *claimed* as violated. To find these cases, we can look through all past cases and pick out those where the same Articles have been claimed. Because all cases in our corpus deal with the same subset of law, by the nature of precedent, all past decisions are all binding with respect to future ones. We call this the **claimed precedent**. However, in practice, considering all claimed precedent cases yields a large number of precedent cases. Thus, we also consider a more narrow view. Specifically, we focus on those cases the judge cites when drafting their judgment. We call this the **cited precedent**. Each approach has its limitations. Under the narrow view (cited precedent), we have a non-exhaustive list of cases that only show us the reasoning of a *single* judge.⁵ Under the broader view (claimed precedent), we consider many cases that are not directly relevant to the facts of the case at hand. The claimed–cited distinction plays a crucial role in our analysis. Correlations are broken down into cited and claimed precedent in Tab. 2

Per-case vs. Per-article. A lawyer typically claims multiple Articles as violated, and, likewise, a judge may find that multiple Articles were breached. Thus, when selecting precedent cases, on one hand, we can consider those where all the claimed Articles are shared. Or, on the other hand, we can consider a piecemeal approach where we select all cases that share *at least* one Article. We

⁵A judgement may be given by a tribunal of judges, each giving a different argument. In our dataset, the reasoning towards the decision is reported as a single argument.

call the former approach the **per-article approach** and the latter the **per-case approach**. The advantage of the per-case approach is that it selects for highly related cases because the cases have an identical set of claimed Articles. The advantage of the per-article approach is that it allows for a more detailed analysis—a case might be an important precedent with respect to a judge’s decision on only a single Article. While we report our main results with per-case correlations, we include an analysis of per-article correlations in §8.2.

4.4 Recovering the Taxonomy

Conveniently, the above taxonomy can be recovered post hoc from the case transcripts. All we need is to construct a citation network. We can then identify the relationship between a cited case and the citing case as follows. First, we identify the outcomes of both the citing case and the cited case. If the positive outcomes of the cited case are the same as those of the citing case, there is a relationship of **applied-positive** precedent. If the negative outcomes are the same, the relationship is that of **applied-negative** precedent. To extract cases that constitute a distinguished precedent, we first identify if the case at hand and the cited case share the same claims. If they do, we move forward to distinguish between positive and negative outcomes. When the positive outcomes of the cited case are the same as the negative outcomes of the citing case, we label the relationship as **distinguished-positive**. Conversely, if the negative outcomes of the cited case are the same as the positive outcomes of the citing case, we label the relationship as **distinguished-negative**.

Note that because we consider positive and negative precedent separately, the relationship between any two cases can simultaneously be both positive and negative, though it can not be both applied and distinguished. Some cited cases might not fall under any of the above categories. We include an example of how we process cases into the taxonomy in App. B.

5 An Improved ECtHR Corpus

We modify the ECtHR corpus, a popular legal NLP dataset, for our experiments. Specifically, we use the scrape of the corpus from Valvoda et al. (2021). An ECtHR case can be separated into two sections: facts and arguments. The facts of the case, denoted $f \in \Sigma^*$, contain text describing the events that had brought the claimant to court. This

section outlines the events that might give rise to a particular breach of law. Arguments of the case on the other hand contain a transcription of the reasoning the judge(s) employed in reaching the conclusion over whether or not a breach of law has actually occurred. The argument section is where the judges cite previous cases, i.e., the precedent.

We first extract the necessary citations to previous cases. Specifically, for every case in our dataset, we extract the names of all the cases cited in the arguments section. We use regular expressions for this purpose because citations in the ECtHR corpus follow a regular form.⁶ The first author manually inspected the extracted citations to ensure a high-quality citation network. Then, we remove all the cases in the training set of our dataset that have *not* been cited. This step ensures that every training case is a precedent of at least one test case. Then we remove all the cases in the test set of our corpus which do not cite any precedent. We do this to ensure that, for each case in our test set of the corpus, there is a precedent cited from our training corpus. Details about the number of cases per Article and the split between the training, validation, and test set are contained in App. A.

6 Neural Models of Legal Outcome

We consider two different instantiations of Eq. (1), each of which corresponds to a different legal outcome prediction approach from the literature.

Simple model. Under Chalkidis et al.’s (2022) model, legal outcome prediction is treated as a binary classification task with the outcome of a case simplistically assumed to be either + or an amalgamation of – and \emptyset . Using a pre-trained language model as an encoder enc , we define the following model:

$$\mathbf{h} = \text{enc}(\mathbf{f}) \quad (7a)$$

$$\mathbf{e} = W_k^{(1)} \text{ReLU}(W^{(2)} \mathbf{h}) \quad (7b)$$

$$p(O_k = + | \mathbf{f}) = \text{sigmoid}(\mathbf{e}_k)_+ \quad (7c)$$

where the multi-layer perceptron (MLP) is parametrized by $W_k^{(1)} \in \mathbb{R}^{1 \times d_2}$ and $W^{(2)} \in \mathbb{R}^{d_2 \times d_1}$. The input $\mathbf{h} \in \mathbb{R}^{d_1}$ is a high-dimensional representation of the facts computed by the encoder.

⁶ECtHR has issued a note on the rules on case citations [here](#). In the ECtHR corpus, the earliest cases are from the year 2001 and the relevant citation styles they follow are the two used between 1999-2015 and from 2016 onwards. Both follow a regular format, allowing us to generate our citation dataset.

Joint model. We also implement Valvoda et al.’s (2023a) model, which treats outcome prediction as a three-way classification task:

$$\mathbf{h} = \text{enc}(\mathbf{f}) \quad (8a)$$

$$\mathbf{e} = W_k^{(3)} \text{ReLU}(W^{(4)} \mathbf{h}) \quad (8b)$$

$$p(O_k = o_k | \mathbf{f}) = \text{softmax}(\mathbf{e}_k)_{o_k} \quad (8c)$$

where $W_k^{(3)} \in \mathbb{R}^{3 \times d_2}$ and $W^{(4)} \in \mathbb{R}^{d_2 \times d_1}$ are the parameters of the MLP. The advantage of this approach is that it directly models both positive and negative outcomes, which are closely related to positive and negative precedent. Thus, we expect the Joint model to be better at utilizing them.

Choice of enc. For both models, we consider two choices for the encoding $\text{enc} : \Sigma^* \rightarrow \mathbb{R}^{d_1}$: (1) BERT (Devlin et al., 2019) and (2) LEGAL-BERT (Chalkidis et al., 2020). The full experimental details, e.g., hyperparameters, are given in App. D.

7 Results

We report our results in terms of Spearman’s ρ . Overall, we find a weak positive correlation between judges’ precedent and influence scores—the best correlation we find is 0.180. However, due to the fine-grained taxonomy, the results vary between different types of precedent. Inspecting the main results, contained in Tab. 2, we make the following six findings.

Finding #1. We observe that influence scores correlate more strongly with applied precedent than distinguished precedent. In fact, the distinguished precedent is negatively correlated with the influence scores in the case of the Simple model. Under the best Joint model, the highest correlation we find between the distinguished precedent and influence scores is nearly $4 \times$ lower (0.020) when compared to the highest correlation we obtain for the applied precedent (0.079). Therefore, we conclude that current legal outcome prediction models struggle with distinguished precedent.

Finding #2. We find negative precedent correlates less strongly with influence scores than positive precedent. In the case of applied precedent, we observe a correlation with positive precedent of 0.180, which is more than $6 \times$ larger than the correlation we observe with negative precedent (0.027). Therefore, we conclude the models struggle to utilize negative precedent in comparison.

Type	Language Model	F_1		Applied		Distinguished		Overall
				Cited	Claimed	Cited	Claimed	
Simple	BERT	0.64	Positive	0.013	0.148	-0.002	-0.012	0.137
			Negative	0.004	0.037	-0.002	-0.023	
Joint	LEGAL-BERT	0.67	Positive	0.013	0.180	-0.003	-0.026	0.157
			Negative	0.003	0.027	-0.002	-0.023	
	BERT	0.66	Positive	0.008	0.079	0.001	0.020	0.084
			Negative	0.003	0.023	0.001	0.002	
LEGAL-BERT	0.68	Positive	0.008	0.078	0.000	0.020	0.085	
		Negative	0.004	0.031	0.000	0.004		

Table 2: We report per-case Spearman’s ρ between influence scores and our taxonomy of precedent, F_1 scores for the models under consideration and the *overall* correlation between influence scores and any type of precedent.

Finding #3. When analyzed through the lens of influence scores, neither model correctly emulates the full spectrum of precedential reasoning. While the influence scores under the Joint model correlate positively with every type of precedent, we observe an overall correlation of 0.085, which is weaker than the overall correlation of 0.157 that we observe with the Simple model’s influence scores. Because the Simple model is negatively correlated with distinguished precedent, neither of the models under consideration seems to be able to robustly reason with the full set of precedential operations.

Finding #4. We observe that pre-training on legal data is beneficial for improving precedent correlations. The LEGAL-BERT-based models exhibit higher correlations with precedent when compared to the BERT-based models. Under the Simple model, in particular, we see an increase from a correlation of 0.137 to a correlation of 0.157 when we use LEGAL-BERT.

Finding #5. When we compare the correlation for cases a judge has cited with the correlation for claimed precedent cases, we observe that the latter approach magnifies the results. In the case of a positive correlation, the correlations become more positive. In the case of a negative correlation, the correlations become more negative. The effect is considerable—in many instances, the difference is more than an order of magnitude.

Finding #6. We find that higher F_1 scores do not directly translate to more human-like legal reasoning. Where the Joint models outperform the Simple models in terms of F_1 , the overall ρ is higher for the Simple models. Overall, the worst model in terms of F_1 (Simple BERT: 0.64

F_1) exhibits a correlation of 0.14 with precedent. Whereas the best model in terms of F_1 (Joint LEGAL-BERT: 0.68 F_1) exhibits a correlation of merely 0.085 with the precedent.

Conclusion. In summary, we find that current models do well at encoding one particular type of precedent, applied-positive precedent, and not so well at the other types of precedent. We believe the culprit is the lack of sophisticated legal reasoning. To distinguish a case is to identify how two seemingly similar cases differ. This requires in-depth legal knowledge. Our work indicates that this level of legal understanding is beyond the capabilities of the current neural models of legal outcome prediction models.

8 Further Analysis

We extend our analysis with the following three studies to better understand the models’ behavior. First, we analyze the precedent from the perspective of the model. Then, we analyze the model behavior per individual legal Article. Finally, in App. C, we develop a novel metric, which confirms the trends we report in §7 from the perspective of precedent retrieval.

8.1 Model-based Precedent

Current legal outcome predictors often fail to predict the correct outcomes. Therefore, the weak correlations reported above might simply be due to the model relying on the precedent which supports the outcome the model predicts, rather than the outcome a human judge has reached. To investigate this possibility, we develop two ways of selecting precedent in line with the models’ prediction of the




Figure 1: Per-article Spearman correlation between influence functions and cited negative-distinguished precedent. The Simple model is on the left and the Joint model is on the right; both models are LEGAL-BERT based.

Type	LM	Correct	Model
Simple	BERT	0.149	0.007
	LEGAL-BERT	0.183	0.032
Joint	BERT	0.138	0.010
	LEGAL-BERT	0.148	0.011

Table 3: Spearman’s correlations for the correctly predicted cases and model-based precedent. The correlations are based on *overall* types of precedent.

outcome. First, we compute correlations for only those instances where the model has **correctly predicted** the outcome. Second, we select the precedent cases by taking the model outcome prediction as the ground truth. We call this approach **model-based**. Under both approaches, we study the same types of correlations as in §7. The overall Spearman correlations are reported in Tab. 3.

We find that the correlation with precedent is higher for correctly predicted cases. The best overall correlation for these instances improves from 0.157 to 0.183, a 17% improvement. On the other hand, the model-based results reveal a decrease in correlation with precedent. The best model under this paradigm achieves a correlation of only 0.011, more than an order of magnitude worse than the best model in Tab. 2. This suggests that the models agree with a human judge’s choice of precedent more than with the precedent that would be consistent with the model’s own predictions. This finding casts further doubt on the alignment of the models’ reasoning with the doctrine of precedent. We report full results in App. G.

8.2 Per-article Analysis

So far, we have considered only per-case correlations. We now turn to the per-article results. We report our full per-article results in App. E. In

general, we observe the same trends as those set out in §7. Nonetheless, we find interesting results for the negative distinguished precedent; see Fig. 9. We discuss these results in more detail below.

Under the Simple model, we observe that those Articles that are more heavily represented in the training data exhibit a lower correlation with negative distinguished precedent in comparison with those that are less heavily represented. For example, Articles 3, 6, 8, 10, 13 and P1-1, each of which has a large number of examples in the training set, correlate more negatively with negative distinguished precedent than their less well-represented counterparts. Therefore, we contend that it is unlikely that introducing additional training data will solve the problem of limited precedential reasoning in neural models of legal outcome. In fact, scaling up the data seems to result in less human-like behavior. Instead, our results suggest that we need to develop new models with a better legal inductive bias to address this issue.

When it comes to distinguished precedent, Joint model exhibits some aspects of such inductive bias. For most Articles, and Articles 8 and 13 in particular, the model seems to utilize distinguished precedent. Nonetheless, for Article 6, which is the most represented Article in the dataset, we still observe a negative correlation. While Joint model appears to be a step in the right direction in some respects, it is still far from using the precedent in a way a human judge would.

9 Related Work

In the wider NLP context, there is a considerable amount of research aimed at developing methods for unraveling what LMs know about language (Alain and Bengio, 2016; Shi et al., 2016; Ettinger et al., 2016; Bisazza and Tump, 2018; Liu et al., 2019; Pimentel et al., 2020). Sub-fields of

NLP dealing with sensitive information, such as biomedical NLP (Garcia-Olano et al., 2021; Jha et al., 2018; Moradi and Samwald, 2021; Sarker et al., 2019; Mullenbach et al., 2018), have set out to build their own explainable AI models to address the concerns about applying a black-box model to sensitive data.

In the past, legal AI researchers working on symbolic methods built explainable models of legal reasoning (Ashley, 1988; Alevn and Ashley, 1997; Risland and Skalak, 1991; Branting, 1991). However, the most recent crop of models utilizing modern NLP methods are inherently opaque. Regardless of whether the domain is Chinese (Feng et al., 2022), Japanese (Yamada et al., 2023), Indian (Malik et al., 2021) or UK law (Valvoda et al., 2018), the machine learning models underlying the work in the area are not easily interpretable.

To address this shortcoming, straightforwardly adopting existing methodology from wider NLP is insufficient. After all, how well a model knows syntax can only go so far in explaining its legal decisions. Using more general methods, such as attention maps, is not particularly viable either. For instance, using attention to find where the model *looks* at legal facts does not provide a legally salient answer as to what the decision is based on (Mullenbach et al., 2018).

Setting aside the question of whether attention can be used for explainability in the first place (Serrano and Smith, 2019; Bastings and Filippova, 2020; Meister et al., 2021; Bibal et al., 2022), models of legal outcome prediction are trained to make decisions based on the facts of a case alone. While the facts of the case are important, inspecting the attention distribution over them can never provide a sufficient explanation. This is because lawyers and judges do not argue by highlighting which facts played a role in their decision. Instead, in their arguments, they justify their decision by relating the facts at hand to previous case decisions. This makes our work different from prior research in two aspects. (1) Where the existing Legal AI work focuses on providing explanations by identifying relevant tokens (Xu et al., 2023), sentences (Malik et al., 2021) or paragraphs (Chalkidis et al., 2021) of the facts of the case, we find the entire relevant cases. (2) Instead of looking inside the input to the model, i.e., at the facts, we look at the training data of the model.

10 Conclusion

We introduce a method based on influence functions for legal explainability. We created a taxonomy of precedent, used it to annotate the ECtHR dataset and ran a series of experiments to identify what precedential operations might be learned, or not, by outcome prediction models. We find that the models we train rely on one type of precedent in particular: applied precedent. This makes them very unlike human judges, who employ a full spectrum of precedential reasoning to justify the outcome of a case. To construct useful models of legal outcomes, we need to develop legally faithful models of law. We hope that our methodology can help in the search for such models.

Indeed, we see two concrete avenues toward legal NLP models that could better align with human precedential reasoning. (1) We can focus on creating more representative datasets of the tasks we would like to model. The facts in current datasets contain superficial indicators of the outcome—they are written as post-hoc justification. Using datasets built around legal briefs, instead of cases, could be a step towards mitigating this issue (Valvoda et al., 2023b). (2) We can better incorporate the court processes in our neural models. To induce precedential reasoning, a natural step would be to incorporate precedent retrieval as part of the model architecture. Much like a judge, the model should be able to explicitly select the precedent that its decision is based on.

Limitations

There are a number of limitations to our work which we would like to acknowledge. First, our work is limited to a study of the English version of the ECtHR corpus. Future work should investigate the French translations of the caselaw, which is available for all cases in the corpus. Additionally, one could experiment with the translations available in other languages, since every case is also translated into the official language of the defendant. Second, there are potential limitations to using our methodology in different jurisdictions. We report the details of our dataset and experiments in App. A and App. D, respectively. Training of the models requires a sufficiently large dataset which could be a limiting factor for smaller jurisdictions with fewer training cases available.

Ethics Statement

The models and methods in this paper are intended purely for scientific research. While we believe that one day legal AI could be a valuable tool in speeding up the legal process, the current state of the research, including ours, is far from being ready for real-world deployment. We would therefore caution anyone from using our work towards automating away the job of a human judge or lawyer. Our results in fact suggest that the current models behave unlike human judges when it comes to precedent. This makes them highly unsuitable for real-world deployment. See our work on legal ethics for a more detailed critique of the current state of Legal AI (Valvoda et al., 2023b).

Acknowledgments

This research is funded by the Nordic Programme for Interdisciplinary Research Grant 105178 and the Danish National Research Foundation Grant no. DNRF169. We would also like to thank Professor Ken Satoh for the fruitful discussions and feedback.

References

- Guillaume Alain and Yoshua Bengio. 2016. [Understanding intermediate layers using linear classifier probes](#). *arXiv preprint arXiv:1610.01644*.
- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoŕiuc-Pietro, and Vasileios Lamos. 2016. [Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective](#). *PeerJ Computer Science*, 2:e93.
- Vincent Aleven and Kevin Ashley. 1997. [Teaching case-based argumentation through a model and examples empirical evaluation of an intelligent learning environment](#). In *Artificial intelligence in education*, volume 39, pages 87–94. IOS Press.
- Kevin Ashley. 1988. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. Ph.D. thesis, USA. Order No: GAX88-13198.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online. Association for Computational Linguistics.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François, and Patrick Watrin. 2022. [Is attention explanation? an introduction to the debate](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland. Association for Computational Linguistics.
- Arianna Bisazza and Clara Tump. 2018. [The lazy encoder: A fine-grained analysis of the role of morphology in neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2871–2876, Brussels, Belgium. Association for Computational Linguistics.
- Henry Black. 2019. *Black’s Law Dictionary*, 11th edition. Thomson Reuters.
- K. Branting, B. Weiss, B. Brown, C. Pfeifer, A. Chakraborty, L. Ferro, M. Pfaff, and A. Yeh. 2019. [Semi-supervised methods for explainable legal prediction](#). In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, ICAIL ’19*, page 22–31, New York, NY, USA. Association for Computing Machinery.
- Karl Branting. 1991. [Reasoning with portions of precedents](#). In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 145–154. ACM.
- 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.
- 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 Linguistics: Human Language Technologies*, pages 226–241, Online. Association for Computational Linguistics.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [LexGLUE: A benchmark dataset for legal language understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume*

- I: Long Papers*), pages 4310–4330, Dublin, Ireland. Association for Computational Linguistics.
- Ilias Chalkidis and Anders Søgaard. 2022. [Improved multi-label classification under temporal concept drift: Rethinking group-robust algorithms in a label-wise setting](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2441–2454, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Neil Duxbury. 2008. *Distinguishing, overruling and the problem of self-reference*. Cambridge University Press.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. [Probing for semantic evidence of composition by means of simple classification tasks](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 134–139, Berlin, Germany. Association for Computational Linguistics.
- Yi Feng, Chuanyi Li, and Vincent Ng. 2022. [Legal judgment prediction via event extraction with constraints](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 648–664, Dublin, Ireland. Association for Computational Linguistics.
- Diego Garcia-Olano, Yasumasa Onoe, Ioana Baldini, Joydeep Ghosh, Byron Wallace, and Kush Varshney. 2021. [Biomedical interpretable entity representations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3547–3561, Online. Association for Computational Linguistics.
- Ivan Habernal, Daniel Faber, Nicola Recchia, Sebastian Bretthauer, Iryna Gurevych, Indra Spiecker genannt Döhmann, and Christoph Burchard. 2023. [Mining legal arguments in court decisions](#). *Artificial Intelligence and Law*.
- Philipp Hacker, Ralf Krestel, Stefan Grundmann, and Felix Naumann. 2020. [Explainable AI under contract and tort law: Legal incentives and technical challenges](#). *SSRN Electronic Journal*.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kishlay Jha, Yaqing Wang, Guangxu Xun, and Aidong Zhang. 2018. [Interpretable word embeddings for medical domain](#). In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 1061–1066.
- Mi-Young Kim, Juliano Rabelo, Randy Goebel, Masaharu Yoshioka, Yoshinobu Kano, and Ken Satoh. 2023. [COLIEE 2022 summary: Methods for legal document retrieval and entailment](#). In *New Frontiers in Artificial Intelligence: JSAI-IsAI 2022 Workshop, JURISIN 2022, and JSAI 2022 International Session, Kyoto, Japan, June 12–17, 2022, Revised Selected Papers*, page 51–67, Berlin, Heidelberg. Springer-Verlag.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML 17*, page 1885–1894. JMLR.org.
- Grant Lamond. 2016. [Precedent and Analogy in Legal Reasoning](#). In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, spring 2016 edition. Metaphysics Research Lab, Stanford University.
- Reed C. Lawlor. 1963. [What computers can do: Analysis and prediction of judicial decisions](#). *American Bar Association Journal*, 49(4):337–344.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). 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 1073–1094, Minneapolis, Minnesota. 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.
- 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.
- Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. 2021. [Is sparse attention more interpretable?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 122–129, Online. Association for Computational Linguistics.

- Milad Moradi and Matthias Samwald. 2021. [Explaining black-box models for biomedical text classification](#). *IEEE Journal of Biomedical and Health Informatics*, 25(8):3112–3120.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111, New Orleans, Louisiana. Association for Computational Linguistics.
- 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, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Milda Norkute, Nadja Herger, Leszek Michalak, Andrew Mulder, and Sally Gao. 2021. [Towards explainable AI: Assessing the usefulness and impact of added explainability features in legal document summarization](#). In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, CHI EA '21*, New York, NY, USA. Association for Computing Machinery.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Tiago Pimentel, Josef Valvoda, Niklas Stoehr, and Ryan Cotterell. 2022. [Attentional probe: Estimating a module’s functional potential](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11459–11472, Abu Dhabi, United Arab Emirates. 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.
- Edwina Rissland and David Skalak. 1991. [Cabaret: rule interpretation in a hybrid architecture](#). *International journal of man-machine studies*, 34(6):839–887.
- W. Rudin. 1976. *Principles of Mathematical Analysis*. International series in pure and applied mathematics. McGraw-Hill.
- Abeed Sarker, Ari Z. Klein, Janet Mee, Polina Harik, and Graciela Gonzalez-Hernandez. 2019. [An interpretable natural language processing system for written medical examination assessment](#). *Journal of Biomedical Informatics*, 98:103268.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. [Does string-based neural MT learn source syntax?](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1526–1534, Austin, Texas. Association for Computational Linguistics.
- 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, ICISS 2020*, page 204–209, New York, NY, USA. Association for Computing Machinery.
- 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, Varna, Bulgaria. INCOMA Ltd.
- Santosh T.y.s.s, Oana Ichim, and Matthias Grabmair. 2023. [Zero-shot transfer of article-aware legal outcome classification for European court of human rights cases](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 605–617, Dubrovnik, Croatia. Association for Computational Linguistics.
- Josef Valvoda, Ryan Cotterell, and Simone Teufel. 2023a. [On the role of negative precedent in legal outcome prediction](#). *Transactions of the Association for Computational Linguistics*, 11:34–48.
- Josef Valvoda, Tiago Pimentel, Niklas Stoehr, Ryan Cotterell, and Simone Teufel. 2021. [What about the precedent: An information-theoretic analysis of common law](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2275–2288, Online. Association for Computational Linguistics.

Josef Valvoda, Oliver Ray, and Ken Satoh. 2018. [Using agreement statements to identify majority opinion in UKHL case law](#). In *Legal Knowledge and Information Systems - JURIX 2018: The Thirty-first Annual Conference, Groningen, The Netherlands, 12-14 December 2018*, volume 313 of *Frontiers in Artificial Intelligence and Applications*, pages 141–150. IOS Press.

Josef Valvoda, Alec Thompson, Ryan Cotterell, and Simone Teufel. 2023b. [The ethics of automating legal actors](#). *arXiv:2312.00584*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Shanshan Xu, Santosh T.y.s.s, Oana Ichim, Isabella Risini, Barbara Plank, and Matthias Grabmair. 2023. [From dissonance to insights: Dissecting disagreements in rationale construction for case outcome classification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9576, Singapore. Association for Computational Linguistics.

Hiroaki Yamada, Takenobu Tokunaga, Ryutaro Ohara, Akira Tokutsu, Keisuke Takeshita, and Mihoko Sumida. 2023. [Japanese tort-case dataset for rationale-supported legal judgment prediction](#). *arXiv:2312.00480*.

A Dataset Details

We report the breakdown of the dataset in Tab. 4 below. As you can see, the dataset is highly imbalanced, which is why we conduct our per Article analysis in §8.

B An Example of Our Taxonomy

To better understand how we taxonomize precedent, let us consider an imaginary case \mathbf{o} with five precedent candidate cases $\mathbf{o}'_1, \dots, \mathbf{o}'_5$. For the sake of simplicity, we will set these cases in an imaginary legal system, which only has four laws (i.e., $K = 4$). Below, we illustrate how we would assign the type of precedent to each candidate case:

- $\mathbf{o} = [\emptyset, +, -, \emptyset]$ (Case at hand)
- $\mathbf{o}'_1 = [\emptyset, +, \emptyset, \emptyset]$ (Applied positive)
- $\mathbf{o}'_2 = [\emptyset, \emptyset, -, \emptyset]$ (Applied negative)
- $\mathbf{o}'_3 = [\emptyset, \emptyset, +, \emptyset]$ (Distinguished positive)
- $\mathbf{o}'_4 = [\emptyset, -, \emptyset, \emptyset]$ (Distinguished negative)
- $\mathbf{o}'_5 = [\emptyset, \emptyset, \emptyset, +]$ (No precedent)

	Train	Validation	Test
No. of Cases	8344	904	913
Art 2	463	56	55
Art 3	1223	169	173
Art 5	1255	176	154
Art 6	4434	276	272
Art 7	24	2	3
Art 8	638	76	111
Art 9	37	4	4
Art 10	268	39	66
Art 11	102	29	33
Art 13	1124	77	71
Art 14	119	17	11
P1-1	1307	130	116
P1-3	43	4	1
P4-2	31	3	4

Table 4: The number of cases in the precedent corpus and the representation of each ECtHR Article.

C A New Evaluation Metric

We now develop a novel evaluation metric to determine the degree to which a case’s influence corresponds to whether or not it was used as a precedent by the model. To map a case’s influence to the binary decision of whether or not the case was used as precedent, we design a simple classifier. And, to deal with the imbalanced nature of the classification (there are far more non-relevant cases than relevant ones), we evaluate the classifier with F_1 . We omit this analysis from the main text because the conclusions drawn from it are nearly identical to those made in the main with more standard methods.

We now define the classifier formally. Given an influence score $\iota(\mathbf{z}, \mathbf{z}') = x$, we construct the classifier:

$$p(C_n = 1) = \frac{\exp(ax_n + b)}{1 + \exp(ax_n + b)} \quad (9)$$

where $C_n = 1$ indicates the n^{th} training case is considered precedent. The function to the right of Eq. (9) is a sigmoid function. Thus, we know that the classification boundary induced by Eq. (9) is:

$$ax + b \geq 0 \implies x + \frac{b}{a} \geq 0 \implies x \geq -\frac{b}{a} \quad (10)$$

Therefore, under Eq. (9), any parameters a and b give us a real value $-\frac{b}{a}$ such that a case is predicted to be precedent iff its influence is greater than $-\frac{b}{a}$.

		Model	Rand.	Gain
Simple	BERT	0.25	0.18	+0.07
	L-BERT	0.27	0.18	+0.09
Joint	BERT	0.21	0.18	+0.03
	L-BERT	0.21	0.18	+0.03

Table 5: The F_1 scores achieved by our models on the precedent retrieval task on the ECtHR corpus.

To learn the two parameters $a = \theta_0, b = \theta_1$, we optimize the log-likelihood:

$$\sum_{n=1}^N \log p(C_n = 1) + \kappa \|\boldsymbol{\theta}\|_2^2 \quad (11)$$

where $\kappa > 0$ is a small positive number used to break ties.

Precedent Classifier Results. We report the results of the above classifier in Tab. 5. In comparison to the random baseline, the influence scores are consistently a better predictor of whether or not a case has been used as a precedent. This finding holds over all models under consideration. The results also confirm our previous findings—Valvoda et al.’s (2023a) model underperforms Chalkidis and Søgaard’s (2022) model. Furthermore, we find that a classifier on top of a model’s influence scores does not score highly on the precedent prediction task. To contextualize these results, however, we note the results of Task 1 of the past two years of COLIEE competition, which is similar to our precedent retrieval above (but over a different dataset), indicate that this is a difficult problem to solve even for legal information retrieval models. For context, the best-performing teams in 2021 and 2022 scored 0.19 F_1 and 0.37 F_1 respectively on the data from the Federal Court of Canada case law (Rabelo et al., 2022; Kim et al., 2023).

D Experimental Details

In this section, we provide the details of our experimental setup.

Language Model Encoders. We use two Transformer-based (Vaswani et al., 2017) pre-trained language models for our experiments. Following Chalkidis et al. (2019, 2022) we choose to work with base versions of BERT and LEGAL-BERT architectures.

- We use BERT (Devlin et al., 2019) because it is both a popular language model for benchmarking NLP models and also places at the top of LexGLUE (Chalkidis et al., 2022) leaderboard for Task A—outcome classification, which most closely resembles what our models are trained for. It also performs the best in the recent reformulation of the task by Valvoda et al. (2023a).
- Because BERT is not specifically trained on legal text, we additionally use LEGAL-BERT (Chalkidis et al., 2020), which has been fine-tuned on legal data. This allows us to see if exposure to legal language might affect the use of precedent by the model.

Training Procedure. The models are trained end-to-end by minimizing the cross-entropy between the model and the empirical distribution of the training data. We conduct a light hyper-parameter search over the learning rate $\{3e^{-4}, 3e^{-5}, 3e^{-6}\}$, size of hidden states d_2 in the MLP $\{50, 100, 200, 300\}$ and dropout rate $\{0.1, 0.2, 0.3, 0.4\}$. We train our models for a maximum of ten epochs with early stopping on validation loss.⁷ The models are trained for a maximum of 2 hours on a single GPU. We report our results for the model that achieves the lowest loss on the validation set.

Influence Implementation. We rely on a PyTorch (Paszke et al., 2019) implementation of Koh and Liang’s (2017) method,⁸ a re-implementation of the original TensorFlow code by Pang Wei Koh.⁹

Calculating Correlations. We report our results in terms of Spearman’s ρ . By calculating the influence score of every case in the training corpus on every decision in the test corpus, we can begin to quantify the use of precedent in neural models of legal outcome prediction by studying the correlation of the influence scores with our taxonomy of precedent; see §4. We store the scores as a vector $\mathbf{s} \in \mathbb{R}^M$, where M is the number of train cases times the number of test cases. Then, for each of our four categories of precedent, we create a binary vector $\mathbf{c} \in \{0, 1\}^M$. For each element in \mathbf{s} , \mathbf{c} encodes whether the training case is that particular type of precedent 1 or not 0.

⁷We stop training when validation loss over an epoch stops improving.

⁸A GitHub link to the implementation we rely on.

⁹A GitHub link to the original implementation by Koh.

We compute the correlation between c and s and report Spearman’s ρ . Additionally, we compute the overall correlation over all types of precedent. To do this, we simply assign 1 in s to a case if it falls under *any* of the precedential categories.

E All Per-Article Results

Here we report the results of our per-article analysis. The correlations for individual legal Articles follow the trends reported in §7.

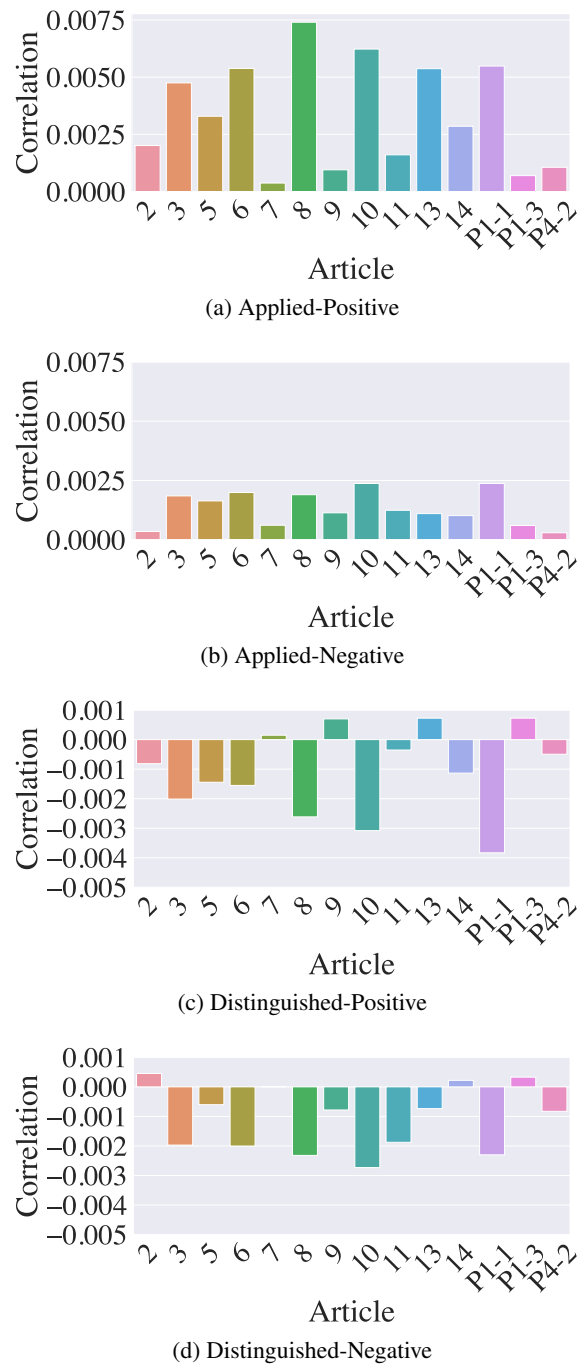


Figure 2: Spearman’s correlation for cited precedent under the Simple model with a BERT encoder..













Figure 3: Spearman's correlation for claimed precedent under the Simple model with a BERT encoder.


Figure 4: Spearman's correlation for cited precedent under the Simple model with a BERT encoder.




(a) Applied-Positive



(b) Applied-Negative




(c) Distinguished-Positive




(d) Distinguished-Negative


Figure 5: Spearman's correlation for all types of precedent under the Simple model with a LEGAL-BERT encoder.




(a) Applied-Positive



(b) Applied-Negative



(c) Distinguished-Positive



(d) Distinguished-Negative

Figure 6: Spearman's correlation for cited precedent under the Joint model with a BERT encoder.







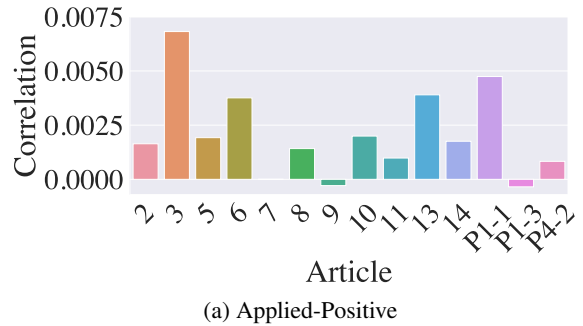





Figure 7: Spearman's correlation for claimed precedent under the Joint model with a BERT encoder.


Figure 8: Spearman's correlation for cited precedent under the Joint model with a LEGAL-BERT encoder.




(a) Applied-Positive



(b) Applied-Negative



(c) Distinguished-Positive



(d) Distinguished-Negative

Figure 9: Spearman's correlation for claimed precedent under the Joint model with a LEGAL-BERT encoder.

F Derivation of the Influence Score

We derive Eq. (5) below. Consider the following bivariate function of θ and ε

$$F(\theta, \varepsilon) \stackrel{\text{def}}{=} \nabla_{\theta} L(\theta) + \alpha \mathcal{R}(\theta) + \varepsilon \ell_z(\theta) = 0 \quad (12)$$

Under suitable regularity conditions on F (Rudin, 1976, Chapter 9), the implicit function theorem guarantees us the following:

1. There exists a function $\theta_z^*(\varepsilon)$ of ε such that $F(\theta_z^*(\varepsilon), \varepsilon) = 0$ in a certain neighborhood;
2. $\theta_z^*(\varepsilon)$ is once differentiable with respect to ε on that neighborhood.

See Rudin (1976, Chapter 9) for more detailed treatment. We now seek to derive $\frac{d\theta_z^*(\varepsilon)}{d\varepsilon}$. We start with the linearity of the gradient

$$0 = \nabla_{\theta} (L(\theta_z^*(\varepsilon)) + \alpha \mathcal{R}(\theta_z^*(\varepsilon)) + \varepsilon \ell_z(\theta_z^*(\varepsilon))) = \nabla_{\theta} L(\theta_z^*(\varepsilon)) + \alpha \nabla_{\theta} \mathcal{R}(\theta_z^*(\varepsilon)) + \varepsilon \nabla_{\theta} \ell_z(\theta_z^*(\varepsilon)) \quad (13)$$

Consider the following manipulation based on taking a derivative with respect to ε :

$$0 = \frac{d0}{d\varepsilon} = \frac{d(\nabla_{\theta} L(\theta_z^*(\varepsilon)) + \alpha \nabla_{\theta} \mathcal{R}(\theta_z^*(\varepsilon)) + \varepsilon \nabla_{\theta} \ell_z(\theta_z^*(\varepsilon)))}{d\varepsilon} \quad (14a)$$

$$= \frac{d\nabla_{\theta} L(\theta_z^*(\varepsilon))}{d\varepsilon} + \alpha \frac{d\nabla_{\theta} \mathcal{R}(\theta_z^*(\varepsilon))}{d\varepsilon} + \frac{d\varepsilon \nabla_{\theta} \ell_z(\theta_z^*(\varepsilon))}{d\varepsilon} \quad (14b)$$

$$= \frac{d\nabla_{\theta} L(\theta_z^*(\varepsilon))}{d\varepsilon} + \alpha \frac{d\nabla_{\theta} \mathcal{R}(\theta_z^*(\varepsilon))}{d\varepsilon} + \nabla_{\theta} \ell_z(\theta_z^*(\varepsilon)) + \varepsilon \frac{d\nabla_{\theta} \ell_z(\theta_z^*(\varepsilon))}{d\varepsilon} \quad (14c)$$

$$= \nabla_{\theta}^2 L(\theta_z^*(\varepsilon)) \frac{d\theta_z^*(\varepsilon)}{d\varepsilon} + \alpha \nabla_{\theta}^2 \mathcal{R}(\theta_z^*(\varepsilon)) \frac{d\theta_z^*(\varepsilon)}{d\varepsilon} + \nabla_{\theta} \ell_z(\theta_z^*(\varepsilon)) + \varepsilon \nabla_{\theta}^2 \ell_z(\theta_z^*(\varepsilon)) \frac{d\theta_z^*(\varepsilon)}{d\varepsilon} \quad (14d)$$

Now, we have

$$-\nabla_{\theta} \ell_z(\theta_z^*(\varepsilon)) = \nabla_{\theta}^2 L(\theta_z^*(\varepsilon)) \frac{d\theta}{d\varepsilon} + \alpha \nabla_{\theta}^2 \mathcal{R}(\theta_z^*(\varepsilon)) \frac{d\theta}{d\varepsilon} + \varepsilon \nabla_{\theta}^2 \ell_z(\theta_z^*(\varepsilon)) \frac{d\theta}{d\varepsilon} \quad (15a)$$

$$= (\nabla_{\theta}^2 L(\theta_z^*(\varepsilon)) + \alpha \nabla_{\theta}^2 \mathcal{R}(\theta_z^*(\varepsilon)) + \varepsilon \nabla_{\theta}^2 \ell_z(\theta_z^*(\varepsilon))) \frac{d\theta}{d\varepsilon} \quad (15b)$$

Evaluating the expression at $\varepsilon = 0$, we arrive at

$$\left. \frac{d\theta_z^*(\varepsilon)}{d\varepsilon} \right|_{\varepsilon=0} = -(\nabla_{\theta}^2 L(\theta^*) + \alpha \nabla_{\theta}^2 \mathcal{R}(\theta^*))^{-1} \nabla_{\theta} \ell_z(\theta^*) \quad (16)$$

where θ^* are the parameters obtained in Eq. (3).

G Additional Results

Here are the full results for our analysis in §8.1.

Type	Language Model	F ₁		Applied		Distinguished		Overall
				Cited	Claimed	Cited	Claimed	
Simple	BERT	0.64	Positive	0.013	0.156	-0.001	0.006	0.149
			Negative	0.004	0.024	-0.002	-0.011	
	LEGAL-BERT	0.67	Positive	0.015	0.205	-0.003	-0.024	0.183
			Negative	0.004	0.025	-0.003	-0.025	
Joint	BERT	0.66	Positive	0.011	0.136	0.002	0.031	0.138
			Negative	0.002	0.003	0.001	0.003	
	LEGAL-BERT	0.68	Positive	0.009	0.146	0.002	0.017	0.148
			Negative	0.001	0.022	0.001	0.009	

Table 6: Spearman correlations between the influence scores and the precedent for cases where the outcome was correctly predicted.

Type	Language Model	F ₁		Applied		Distinguished		Overall
				Cited	Claimed	Cited	Claimed	
Simple	BERT	0.64	Positive	0.002	0.004	0.004	0.042	0.007
			Negative	0.000	-0.027	0.001	0.009	
	LEGAL-BERT	0.67	Positive	0.005	0.029	0.003	0.038	0.032
			Negative	0.000	-0.011	0.000	0.000	
Joint	BERT	0.66	Positive	0.002	0.002	0.003	0.035	0.010
			Negative	0.001	0.004	0.002	0.006	
	LEGAL-BERT	0.68	Positive	0.003	0.002	0.002	0.032	0.011
			Negative	0.000	0.008	0.002	0.013	

Table 7: Spearman correlations between the influence scores and the precedent for cases where the precedent is model-based.