

On the Role of Negative Precedent in Legal Outcome Prediction

Anonymous TAFL submission

Abstract

Every legal case sets a precedent by developing the law in one of the following two ways. It either expands its scope, in which case it sets positive precedent, or it narrows it down, in which case it sets negative precedent. While legal outcome prediction, which is nothing other than the prediction of positive precedents, is an increasingly popular task in AI, we are the first to investigate negative precedent prediction by focusing on negative outcomes. We discover an asymmetry in existing models' ability to predict positive and negative outcomes. Where state-of-the-art outcome prediction models predicts positive outcomes at 75.06 F_1 , they predicts negative outcomes at only 10.09 F_1 , worse than a random baseline. To address this performance gap, we develop two new models inspired by the dynamics of a court process. Our first model significantly improves positive outcome prediction score to 77.15 F_1 and our second model more than doubles the negative outcome prediction performance to 24.01 F_1 . Despite this improvement, shifting focus to negative outcomes reveals that there is still plenty of room to grow when it comes to modelling law.

1 Introduction

The legal system is inherently adversarial. Every case pitches two parties against each other: The *claimant*, who alleges their rights have been breached, and the *defendant*, who denies breaching those rights. For each claim of the claimant, their lawyer will produce an argument for which the defendants lawyer will produce a counter argument. In precedential legal systems, both sides will rely on previous case decisions to support their position (Duxbury, 2008; Lamond, 2016; Black, 2019). The claimant will assert that her circumstances are alike those of previous claimants

whose rights have been breached. The defendant, on the other hand, will allege that the circumstances are in fact more alike those of unsuccessful claimants. The judge decides who is right, and by doing so establishes a new precedent. If it is the claimant who is successful in a particular claim, the precedent expands the law by including the new facts in its scope. If it is the defendant who is successful, the law is contracted by rejection of the new facts from its scope. The expansion or contraction is encoded in the case outcome, we will refer to them as **positive outcome** and **negative outcome**, respectively.

Positive and negative outcomes are equally binding, which means that the same reasons that motivate the research of the positive outcome apply to the negative outcome. Both are important for computational legal analysis, a fact that has been known at least since Lawlor (1963).¹ However, the *de facto* interpretation of precedent in today's legal NLP landscape, studies only the positive outcome, by predicting the breach of law given the facts of a legal case (Aletras et al., 2016). Several researchers have shown that a simple model can achieve a very high performance for the case outcome prediction task (Aletras et al., 2016; Chalkidis et al., 2019; Clavié and Alphonso, 2021; Chalkidis et al., 2021b), a finding that has been replicated for a number of jurisdictions (Zhong et al., 2018; Xu et al., 2020).

In this work, we reformulate outcome prediction as the task of predicting both the positive and negative outcome given the facts of the case. Our results demonstrate that while a simple BERT-based classification model can predict positive outcome with a 75.06 F_1 , it predicts the negative outcome with only 10.09 F_1 , falling short of a random baseline of 11.12 F_1 . This naturally raises the question: What causes such asymmetry? In §8, we argue that this disparity is caused by a

¹He describes them as *pro-precedent* and *con-precedent*.

necessity for deeper legal understanding when it comes to negative outcome prediction.

Searching for a way to better predict negative outcomes, we hypothesise that building a probabilistic model that is more faithful to the legal process will improve both negative and positive outcome prediction. To test this hypothesis we develop two such models. Our first model, which we call the **joint model**, is trained to jointly predict positive and negative outcome. Our second model, which we call the **claim–outcome model**, enforces the relationship between the claims and outcomes. While the joint model significantly outperforms state-of-the-art models on positive outcome prediction with 77.15 F_1 , the claim–outcome model doubles the state-of-the-art models on negative outcome prediction at 24.01 F_1 . We take this as strong evidence that building neural models of the legal process should incorporate domain-specific knowledge of how the legal process works.

2 Judicial Process

In order to motivate our two models of outcome, it is necessary to first understand the formation of law as a global phenomenon. Broadly speaking, the legal process can be understood as a task of narrowing down the legal space where the breach of law might have taken place. Initially, before the legal process begins, the space includes all the law there is, i.e., every legal Article. It is the job of the lawyer to narrow it down to only a small number of Articles, a subset of all law. Finally, the judge determines which of the claimed Articles, if any, has been violated. We can therefore observe two distinct interactions between the real world and the law: first, when a lawyer connects them via a claim, second when a judge connects them via an outcome.

In practice, this means that judges are constrained in their decision. They cannot decide that a law has been breached unless a lawyer has claimed it was breached. This also means that lawyers actively shape the outcome by forcing a judge to consider a particular subset of law. In doing so a lawyer defines the set of laws from which the judge decides on the outcome. The power of a lawyer is also constrained. On one hand, lawyers will want to claim as many legal Articles as possible, on the other there are only so many legal Articles which are relevant to their

clients needs. There are two rules that arise from the interaction of a lawyer and a judge. First, positive outcome is a subset of claims. Second, negative outcome consists of the unsuccessful claims, i.e. the claims the judge rejected.

There is a close relationship between claims and negative outcomes: If we knew the claims the lawyer had made, we could define negative outcome as exactly those Articles that have been claimed but were not found to be violated. However, much like outcomes are a product of judges, claims are a product of lawyers and, unlike facts, they are not known before human legal experts interact with the case. Therefore, to study the relationship of outcomes and facts, one can not simply rely on claims as an additional input to the model. The only input available are the facts.

What’s wrong with current work. Existing models of outcome ignore the role of a legal claim. Under the current formulation of the task, models determine whether a case corresponds to a violation of each Article, i.e., they predict a vector in $\{0, 1\}^K$ where 1 indicates a positive outcome and K is the number of legal Articles under consideration. However, this definition does not help us define when the result is a negative outcome because 0 is ambiguous. In this setting 0 can indicate either “not claimed” or “negative outcome.” Without claims, the existing models implicitly assume that all Articles have been claimed, which is never the case. It is this assumption that allows the existing models to neglect negative outcomes.

Reformulating the task. How should negative outcome then be modelled? Given the domain specific knowledge about the interaction of a judge and a lawyer, our position is that models that predict outcomes should model the claims and outcomes *together*. Towards this, we first need information about which laws have been *claimed*. In §6, we discuss the creation of a new corpus which equips us for this task. In the next section we develop two models that jointly predict outcomes and claims using two basic assumptions about how the law operates. We believe that our reformulation of the task has two advantages. One, considering positive and negative outcomes is a step towards better evaluation of legal outcome prediction models. Two, incorporating the roles of a judge and a lawyer within the models of outcome is a step towards better models of law.

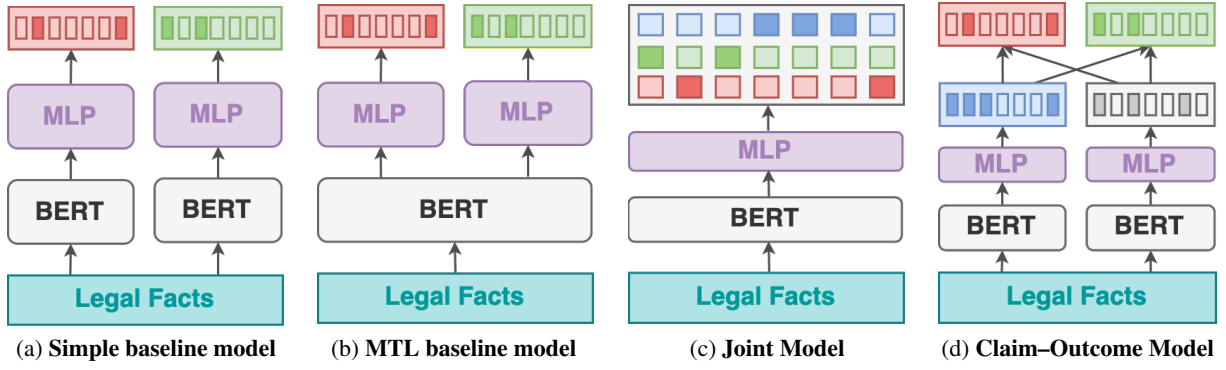


Figure 1: The models under consideration. Green and red boxes represent positive and negative outcomes respectively. Blue boxes represent claims. All our models are trained with facts as the input.

3 Law-Abiding Models

In this section we formulate our two probabilistic models of law. Our law-abiding models are built on top of the two assumptions described below.

Notation. We define a probability distribution over three random variables.

- O is a random variable ranging over $\mathcal{O} = \{+, -, \emptyset\}$, which correspond to positive, negative and null outcome, respectively. The null outcome refers to all the unclaimed Articles, i.e., all Articles which are not positive or negative. The values of O are denoted $o \in \mathcal{O}$. We write o_k for the outcome of the k^{th} Article. \mathbf{O} is a random variable ranging over \mathcal{O}^K where K is the number of Articles we consider. The values of \mathbf{O} are denoted as $\mathbf{o} \in \mathcal{O}^K$.
- The random variable C ranging over $\mathcal{C} = \{Y, N\}$, which corresponds to an Article being claimed or not. The values of C are denoted $c \in \mathcal{C}$. We write c_k for the claim status of the k^{th} Article. The bolded \mathbf{C} is the same as above, i.e., a random variable ranging over \mathcal{C}^K . The values of \mathbf{C} are denoted as $\mathbf{c} \in \mathcal{C}^K$. Note that claims are the inverse of the null outcome.
- The random variable F which ranges over textual descriptions of facts, i.e., Σ^* for a vocabulary Σ . Elements of Σ^* are denoted as f .

3.1 Joint Model

We begin with a simple assumption that given the facts of a case legal Articles are independent:

Assumption 1 (Conditional Independence). *Conditioned on the facts f , the outcome–claim pairs (o_k, c_k) for Article k are conditionally independent of the outcome–claim pairs for other Articles $k' \neq k$.*

This assumption is based in the origin of each Article as an independent Human Right, related by the spirit of ECHR, but otherwise orthogonal in nature. This is indeed how the law operates in general. A law, whether codified in an Article or a product of precedent, encodes a unique right or obligation. In practice this means that a breach of one law does not determine breach of another. For example, a breach of Article 3 of ECHR (the prohibition of torture) does not entail breach of Article 6 (right to a fair trial). Even breaches of law that are closely related, for example libel and slander, do not entail each other and allegation of each must be considered independently.

By Assumption 1, the joint distribution over outcomes and claims decomposes over Articles as

$$p(\mathbf{O} = \mathbf{o}, \mathbf{C} = \mathbf{c} \mid \mathbf{F} = \mathbf{f}) \quad (1)$$

$$= \prod_{k=1}^K p(O_k = o_k, C_k = c_k \mid \mathbf{F} = \mathbf{f})$$

In the remainder of the text we write $\mathbf{F} = \mathbf{f}$ as just f . We also write $O_k = o_k$ and $C_k = c_k$ as o_k and c_k , respectively, when clear from context.

3.2 Claim–Outcome Model

Our second model builds on the first assumption with a second simple assumption:

Assumption 2 (Claims and Outcomes). *Outcomes*

are only determined for claims, i.e.

$$p(O_k = \emptyset \mid C_k = Y, \mathbf{f}) = 0 \quad (2)$$

$$p(O_k = \emptyset \mid C_k = N, \mathbf{f}) = 1 \quad (3)$$

In English, this assumption says that a judge provides an outcome if and only if a claim is made.

By Assumption 2, we have that each distribution over outcome–claim pairs simplifies into the following:

$$p(O_k = o_k, C_k = c_k \mid \mathbf{f}) = \begin{cases} p(o_k \mid Y, \mathbf{f}) p(Y \mid \mathbf{f}) & \text{if } c_k = Y \wedge o_k \neq \emptyset \\ p(N \mid \mathbf{f}) & \text{if } c_k = N \wedge o_k = \emptyset \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Crucially, Assumption 2 allows us to reduce the problem to two independent binary classification problems. First, we train a claim predictor $p(c_k \mid \mathbf{f})$ that predicts whether a lawyer would claim that the k^{th} Article is relevant to the facts \mathbf{f} . Second, we train an outcome predictor that predicts whether the outcome is + or – given that the lawyer has claimed a violation of Article k .

3.3 Neural Parameterization

We consider neural parametrizations for all the distributions discussed above. At heart of all of our models, is a high-dimensional representation $\text{enc} \in \mathbb{R}^{d_1}$ of the facts \mathbf{f} which is the output of a pre-trained language model fine-tuned for our task (see §5). All our language models rely on \mathbf{f} as their sole input. Except where we indicate otherwise, both the language model weights and classifier weights are learned separately for every model presented below.²

Joint Model. First we parameterise the joint model which gives us a joint distribution over all configurations of o_k and c_k for a specific Article k . In principle, there are six such configurations $\{+, -, \emptyset\} \times \{Y, N\}$. However, after we enforce Assumption 2, we are left with three configurations $\{(+, Y), (-, Y), (\emptyset, N)\}$, to which we can assign non-zero probability: This reduces the problem to a 3-way classification, which we parameterise as follows:

$$p(O_k = o_k, C_k = c_k \mid \mathbf{f}) = \text{softmax}(\mathbf{U}_k \rho(\mathbf{V}_k \text{enc}(\mathbf{f})))_{\langle o_k, c_k \rangle} \quad (5)$$

²For the language models, the weights are fine-tuned from a pre-trained model we initialise from the Hugging Face library.

where $\text{softmax}(\mathbf{x})_i = \frac{\exp x_i}{\sum_{i'} \exp x_{i'}}$, ρ is a ReLU activation function defined as $\rho(x) = \max(0, x)$, $\mathbf{U}_k \in \mathbb{R}^{3 \times d_2}$ and $\mathbf{V}_k \in \mathbb{R}^{d_2 \times d_1}$ are per-article learnable parameters. Thus, in total, we have $K(d_2 + d_2 d_1 + d_3 + d_3 d_1)$ parameters for both classifiers excluding those from the encoder enc.

Claim–Outcome Model. We parametrise the claim–outcome model as two binary classification tasks. One predicting the claims, the other predicting positive outcomes. For the latter binary classification task, one class corresponds to +, while the other to both – and \emptyset . This leads to the following pair of binary classifiers:

$$p(C_k = Y \mid \mathbf{f}) = \sigma(\mathbf{u}_k \cdot \rho(\mathbf{V}_k \text{enc}(\mathbf{f}))) \quad (6)$$

$$p(O_k = + \mid C_k = Y, \mathbf{f}) = \sigma(\mathbf{u}'_k \cdot \rho(\mathbf{V}'_k \text{enc}'(\mathbf{f})))$$

where $\mathbf{u}_k \in \mathbb{R}^{d_2}$, $\mathbf{V}_k \in \mathbb{R}^{d_2 \times d_1}$, $\mathbf{u}'_k \in \mathbb{R}^{d_3}$, and $\mathbf{V}'_k \in \mathbb{R}^{d_3 \times d_1}$ are learnable parameters, $\sigma(x) = \frac{1}{1 + \exp(-x)}$ is the sigmoid function, and enc and enc' are two separate encoders. In total, we have $K(d_2 + d_1 d_2 + d_3 + d_1 d_3)$ parameters excluding those from the encoder enc. We use primed symbols to denote separately learned parameters. Given these probabilities, we can marginalise out the claims to obtain the probability of a positive outcome:

$$\begin{aligned} p(O_k = + \mid \mathbf{f}) & \quad (7) \\ &= p(O_k = + \mid C_k = Y, \mathbf{f}) p(C_k = Y \mid \mathbf{f}) \\ & \quad + p(O_k = + \mid C_k = N, \mathbf{f}) p(C_k = N \mid \mathbf{f}) \\ & \stackrel{(1)}{=} p(O_k = + \mid C_k = Y, \mathbf{f}) p(C_k = Y \mid \mathbf{f}) \end{aligned}$$

where (1) is true because $p(O_k = + \mid C_k = N, \mathbf{f})$ is always zero (since by Assumption 2 no positive outcome can be set on an unclaimed case).

We then predict the probability of negative outcome as the inverse of the probability of a positive outcome multiplied by the probability of a claim:

$$p(O_k = - \mid \mathbf{f}) = (1 - p(O_k = + \mid Y, \mathbf{f})) p(C_k = Y \mid \mathbf{f}) \quad (8)$$

This step enforces that the negative outcome probability is always both lower than that of claims and sums up to 1 with positive outcome. Finally, we have the relation that:

$$p(O_k = \emptyset \mid \mathbf{f}) = p(C_k = N \mid \mathbf{f}) \quad (9)$$

To make a decision, we compute the following argmax that marginalises over claims:

$$o_k^* = \operatorname{argmax}_{o_k \in \mathcal{O}} p(O_k = o_k \mid \mathbf{f}) \quad (10)$$

$$= \operatorname{argmax}_{o_k \in \mathcal{O}} \sum_{c_k \in \mathcal{C}} p(O_k = o_k, C_k = c_k \mid \mathbf{f})$$

Training and Fine-tuning. All models in §3.3 are trained by maximizing the log of the joint distribution $p(\mathbf{o}, \mathbf{c} \mid \mathbf{f})$. We are given a dataset of triples $\mathcal{D} = \left\{ \left(\mathbf{o}^{(n)}, \mathbf{c}^{(n)}, \mathbf{f}^{(n)} \right) \right\}_{n=1}^N$. Due to independence assumption made, this additively factorises over articles $\sum_{n=1}^N \log p(\mathbf{o}^{(n)}, \mathbf{c}^{(n)} \mid \mathbf{f}^{(n)}) = \sum_{n=1}^N \sum_{k=1}^K \log p(o_k^{(n)}, c_k^{(n)} \mid \mathbf{f}^{(n)})$. We fine-tune enc jointly for all $p(o_k, c_k \mid \mathbf{f})$.

4 Baselines

We contextualise the performance of the joint and claim–outcome model with a number of baselines. As a starting point we build a simple classification model trained to predict only positive or negative outcome separately, see Fig. 1a. We further want to test whether the advantage of our joint model stems from encoding the relationship between positive and negative outcome, or whether it is down to simply training on more data. We test this by formulating the task as a multi-task learning objective, see Fig. 1b. While this model is trained on the same amount of data as our joint model, it does not explicitly encode the relationship between positive and negative outcomes.

A Simple Baseline. For our simple baseline model we formulate the positive and negative outcome prediction as a multi-label classification task. Despite its being conceptually simple, this model achieves state-of-the-art on the task of case outcome prediction (Chalkidis et al., 2021b), which is an analogous task to our positive outcome prediction task. Given the facts of a case \mathbf{f} we directly model the probability that the outcome is positive as a binary classification problem where the first class is the positive + and the second is the union of negative and unclaimed $\{-, \emptyset\}$. Likewise, we *separately* model the probability that the outcome is negative as a binary classification problem where the first class is the negative – and the second is the union of positive and unclaimed $\{+, \emptyset\}$. To this end, we

define a pair of binary classifiers:

$$p(O_k = + \mid \mathbf{f}) = \sigma(\mathbf{u}_k^1 \cdot \rho(\mathbf{V}_k^1 \operatorname{enc}_1(\mathbf{f}))) \quad (11)$$

$$p(O_k = - \mid \mathbf{f}) = \sigma(\mathbf{u}_k^2 \cdot \rho(\mathbf{V}_k^2 \operatorname{enc}_2(\mathbf{f}))) \quad (12)$$

where $\mathbf{u}_k^1 \in \mathbb{R}^{d_2}$, $\mathbf{V}_k^1 \in \mathbb{R}^{d_2 \times d_1}$, $\mathbf{u}_k^2 \in \mathbb{R}^{d_3}$, and $\mathbf{V}_k^2 \in \mathbb{R}^{d_3 \times d_1}$ are the per-article learnable parameters. Thus, in total, we have $K(d_3 + d_3d_1 + d_2 + d_2d_1)$ parameters excluding those from the fine-tuned encoders enc_1 and enc_2 . The encoders enc_1 and enc_2 represent two *different* fine-tunings of the encoder. Note that this approach does *not* model whether or not an Article is claimed, which stands in contrast to the main models proposed by this work.

MTL Baseline. We also consider a version of the simple baseline where we jointly fine-tune a single encoder. Symbolically, this is written as:

$$p(O_k = + \mid \mathbf{f}) = \sigma(\mathbf{u}_k^1 \cdot \rho(\mathbf{V}_k^1 \operatorname{enc}(\mathbf{f}))) \quad (13)$$

$$p(O_k = - \mid \mathbf{f}) = \sigma(\mathbf{u}_k^2 \cdot \rho(\mathbf{V}_k^2 \operatorname{enc}(\mathbf{f}))) \quad (14)$$

where enc is *shared* between the perceptron heads. Other than the sharing, the the MTL baseline is the same as the simple baseline.

Random Baseline. Finally, we provide a simple random baseline by sampling the outcome vectors from discrete uniform distribution. The random baseline is an average performance over 100 instantiations of this baseline.

5 Experimental Setup

Pre-trained Language Models. We obtain the high-dimensional representation $\operatorname{enc}(\mathbf{f})$ by fine-tuning one of the following pre-trained language models with \mathbf{f} as an input. First, we use BERT because it’s a widely used model in legal AI. Second, we use LEGAL-BERT, because it is purpose trained on legal text which should give it an advantage in our setting. Finally we use the Longformer model. Longformer is built on the same Transformer (Vaswani et al., 2017) architecture as BERT (Devlin et al., 2019) and LEGAL-BERT (Chalkidis et al., 2020), but it can process up to 4,096 tokens. We select this architecture because the facts of legal documents often exceed 512 tokens and a model that can process longer documents could therefore be better suited to our needs.

Chalkidis et al. Corpus			
Outcome	Train	Validation	Test
Positive	8046	835	851
Negative	2259	279	289
Claims	8836	985	991
Outcome Corpus			
Outcome	Train	Validation	Test
Positive	7542	844	925
Negative	4413	498	560
Claims	8372	931	1034

Table 1: Number of cases with at least one positive or negative outcome label in the dataset.

Training Details. All our models are trained with a batch size of 16. We conduct hyperparameter optimisation across learning rate $[3e-4, 3e-5, 3e-6]$, dropout $[0.2, 0.3, 0.4]$ and hidden size of $[50, 100, 200, 300]$. We truncate individual case facts to 512 tokens for BERT and LEGAL-BERT or 4,096 tokens for the Longformer. Our models are implemented using the Pytorch (Paszke et al., 2019) and Huggingface (Wolf et al., 2020) libraries. We use Adam for optimization (Kingma and Ba, 2015) and train all our models on 1 Tesla V100 32GiB GPU’s for a maximum of 1 hour. We train for a maximum of 10 epochs.³ The models are trained on the training set, see Tab. 1. We report the results on the test set for the models that have achieved the lowest loss on the validation set.

6 Legal Corpora

We work with the ECtHR corpus, which contains thousands of instances of case law pertaining to the European Convention of Human rights (ECHR). ECtHR cases contain a written description of case facts, which is our f , and information about claims and outcomes. Since positive outcomes are a subset of all claims, the exclusion set of claims and positive outcomes constitutes the set of negative outcomes.

Chalkidis et al. Corpus. To obtain the golden labels for outcomes and claims we first rely on the Chalkidis et al. (2021a) scrape of the ECHR corpus that contains labeled *alleged violations* and *vi-*

olations. The *violations* are case outcomes, while the *alleged violations* are claims.

Outcome Corpus. However, inspecting Chalkidis et al. Corpus, we find that the *alleged violations* do not contain all of the claims. Alleged violations are only the main claims of the case, a subset of all claims that have been raised. To investigate the full set of claims (and negative outcomes) we process the Chalkidis et al. (2019) scrape of the HUDOC⁴ database and extract the full set of claims using regular expressions.⁵

We conduct all our experiments on both corpora. However, not all of the Articles of ECHR are interesting from the perspective of a legal outcome since not all Articles of ECHR can be claimed by a lawyer. First, ECHR comprises of the convention itself and the protocols to the convention. Out of the 51 Articles of the convention, only Articles 2 to 18 contain the rights and freedoms, whereas the remaining Articles pertain to the court and its operation. We therefore restrict our study to predicting the outcome of the core rights. Furthermore, we remove any Articles that do not appear in the validation and test sets. This leaves us with $K = 17$ and $K = 14$ for the Chalkidis et al. Corpus and Outcome Corpus respectively. Tab. 1 shows the number of cases containing negative outcome vs positive outcome across the training/validation/test splits. The full distribution of Articles over cases in both corpora can be found in App. C.

7 Results

Following Chalkidis et al. (2019), we report all results as F_1 micro-averaged.⁶ The bulk of our results is contained in Tab. 2. First, we compare the positive and negative outcome prediction performance on our outcome corpus and find that while the best simple baseline model achieves 75.06 F_1 on positive outcomes, the same model achieves only 10.09 F_1 on negative outcomes. In fact, the model fails to beat our random baseline of 11.12 F_1 on negative outcomes. The same trend holds over all our model architectures, all the underlying language models and both datasets

⁴See App. B for examples from our dataset or HUDOC for all the ECHR caselaw.

⁵The Outcome Corpus is available at *redacted for anonymity*.

⁶We measure significance using the two tailed paired permutation tests with $p < 0.05$.

³All our code is available at: *redacted for anonymity*

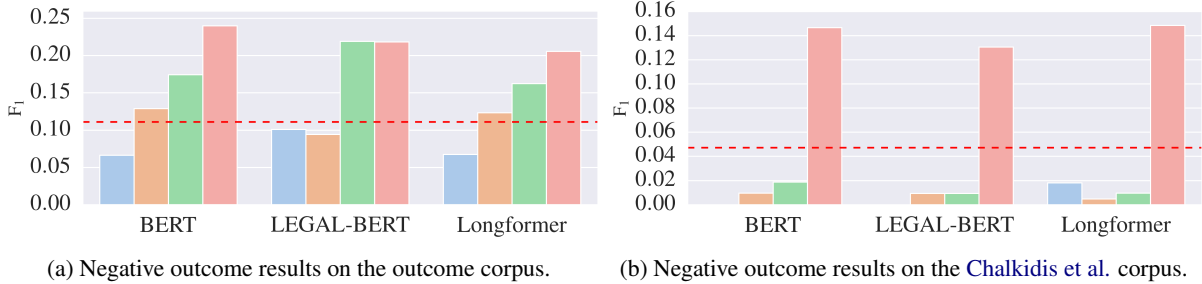


Figure 2: Results for Simple, MTL, Joint and Claim-Outcome models. Dashed line is our random baseline.

Model	Outcome Corpus				Chalkidis et al. Corpus				
	LLM	Pos	Neg	Null	All	Pos	Neg	Null	All
Claim-Outcome	BERT	74.80	24.01	95.53	64.78	63.85	14.65	97.15	58.55
	L-BERT	74.90	21.83	95.49	64.07	64.47	13.05	97.14	58.22
	Longformer	74.23	20.55	95.17	63.32	63.53	14.84	97.21	58.53
Joint	BERT	76.24	17.43	95.46	63.04	65.15	1.87	97.07	54.70
	L-BERT	76.96	21.93	95.71	64.87	67.08	0.94	97.19	55.07
	Longformer	77.15	16.24	95.49	62.96	65.94	0.95	97.11	54.67
MTL Baseline	BERT	75.75	12.90	-	-	63.21	0.95	-	-
	L-BERT	76.73	9.44	-	-	65.00	0.95	-	-
	Longformer	75.83	12.34	-	-	63.36	0.47	-	-
Simple Baseline	BERT	75.06	6.62	-	-	65.04	0.00	-	-
	L-BERT	74.85	10.09	-	-	65.51	0.00	-	-
	Longformer	74.12	6.72	-	-	63.92	1.81	-	-

Table 2: F_1 -micro averaged scores for all the models considered over the two datasets.

under consideration. Every time, the negative outcome performance is significantly lower than that of positive outcomes. Therefore, our first conclusion is that negative outcome is simply harder to predict than its positive counterpart.

Now we turn to the comparison of different model architectures. First, we consider the task of negative outcome prediction. Here, we observe a large and significant improvements using our claim–outcome model, see Fig. 2a and Fig. 2b. Our claim–outcome model is significantly better than every baseline model under consideration, a finding that holds over three underlying language models and both datasets. A single exception to this rule is the joint model, which narrowly beats our claim–outcome model (by 0.1) on the outcome corpus using the LEGAL-BERT LLM. Overall, where the best claim–outcome model achieves 24.01 F_1 on the outcome corpus and 14.84 F_1 on the Chalkidis et al. corpus, the best simple

baseline model only achieves 10.09 and 1.81 F_1 , respectively. Therefore, our second conclusion is that enforcing the relationship between claims and outcomes significantly improves negative outcome prediction. We expand on this in App. A.

Turning to the positive outcome prediction task, the F_1 is comparable between the simple baseline and our claim–outcome model. The joint model on the other hand improves significantly over either baseline and achieves the best F_1 on both the outcome corpus and the Chalkidis et al. corpus (77.15 and 67.08 F_1 respectively). Since the simple baseline model using pre-trained BERT is the state-of-the-art model for positive outcome prediction (Chalkidis et al., 2021b), our third conclusion is that jointly training on positive and negative outcomes is a significantly better way of learning how to predict a positive outcome of a case.

Comparing BERT, LEGAL-BERT and Longformer as the underlying language mod-

els, we find that neither LEGAL-BERT nor Longformer consistently outperform BERT. This is surprising at first sight, since LEGAL-BERT nor Longformer should in theory excel on long legal text. However, these results are consistent with LexGLUE task A (Chalkidis et al., 2021b), which is most similar to our experiments. Moreover, in our setting it seems that LEGAL-BERT has significant positive effect on negative outcome prediction for the joint model and is the underlying language model for our best outcome corpus model. Longformer sets the best positive outcome performance on the same corpus. We therefore find both longer document encoding and legal language pre-training useful in certain settings, although it seems that the choice of model architecture has a larger effect on the performance than the choice of the language model.

The effect varies between corpora and classification model, but for negative precedent prediction, LEGAL-BERT and Longformer are both slightly worse for the precedent corpus (Fig. 2a), but significantly better for the Chalkidis et al. corpus (Fig. 2b) when compared to BERT.

Finally, we turn to the question of what is the best model of outcome; is the the joint model or claim–outcome model? Towards answering this question we take an average F_1 over all three random variables under consideration; the best model of outcome should do well at distinguishing between positive, negative and null outcome. We find that while the joint model has a slight edge over the claim–outcome model on the outcome corpus (by 0.1), the claim–outcome model substantially improves over the joint model on the Chalkidis et al. corpus by 3.48 F_1 . This leads us to believe that claim–outcome model is overall better model of legal outcome prediction. However, both models are valuable in their own right. Where the joint model improves over the state-of-the-art positive outcome prediction models, the claim–outcome model doubles their performance on the negative outcome task.

8 Discussion

The results reported above raise the question of why models severely underperform on negative outcome prediction. The simplest answer could be the amount of training data that is available for each task. We test this hypothesis by comparing the performance on Articles 8 (796 negative vs.

654 positive examples) and 13 (1197 negative vs. 1031 positive examples) of ECHR where in our outcome corpus there is more training data available for the negative outcome than positive outcome. The results, see Fig. 3, prove that even when the model has more training data for negative outcome than the positive outcome, predicting negative outcome is still harder. In particular, for Article 13 the amount of training data is higher than for Article 8, yet the drop in performance between positive and negative outcome prediction is still dramatic. In fact, while the scores achieved by the claim–outcome model are still low, the other models (except our joint model) fail to predict a single negative outcome correctly for Article 13. We therefore believe that the performance drop is likely to be more related with the complexity of the underlying task, rather than the imbalance of the underlying datasets.

To find a better explanation of the performance asymmetry, we now turn in our discussion to the legal perspective. In precedential jurisdictions, of which ECtHR is one (Zupancic, 2016; Lupu and Voeten, 2010; Valvoda et al., 2021), the decisions of a case are binding on future decisions of the court. Two cases with the same facts should therefore arrive at the same outcome. Of course, in reality, the facts are never the same. Rather, cases with similar circumstances will, broadly speaking, lead to similar outcomes. This is achieved by *applying* the precedent. In such cases, the judge will in effect say that the new case is not substantially different from an already existing case and therefore the same outcome will be propagated.

On the other hand, if the previous precedent is not to be followed, the judge needs to *distinguish* the case at hand from the precedent. Distinguishing the case from the precedent, is a more involved task than applying the precedent. It requires identifying what about the new facts sets the new case apart from the previous one. This can of course be done for both cases with positive and negative outcome. Both can be applied or distinguished.⁷ However, since judges deal with claims, each of which comes with an argument built around the precedent that favours the claimant’s viewpoint, we believe that negative outcomes do overwhelmingly rely on distinguishing the case from the precedent. This is evidenced

⁷Overruling is another option, though it is exceptionally rare at the ECtHR (Dzehtsiarou, 2017).

in the yearly reports of the ECtHR (2020), which list cases where the judges decided to distinguish the facts of the case at hand. Distinguishing almost always leads to a negative outcome. We observe the same trend in our ECtHR corpus.

It might therefore be the case that while there is such a thing as a prototypical positive precedent, there is no prototypical negative precedent. This could explain why the simpler architectures struggle to learn to predict it. While a simpler model is ill-suited for the task since it is trained to find a similarity between the negative outcome cases, our claim-outcome model does not assume that negative outcome cases are similar in the first place. Instead, our model assumes similarity between claims. Since claim prediction can be modelled with a high accuracy (Chalkidis et al., 2021a), we can reveal the negative outcome as a disagreement between a judge and a lawyer (i.e. claims and the outcomes).

We can further identify a possible explanation for the baseline model performance by investigating individual cases in Chalkidis et al. Corpus. Consider the case of *Wetjen and Others v. Germany* (Wetjen). This is a case concerned with Article 8: Right to respect for private and family life. On a superficial level, the case is similar to that of *Shuruk v. Switzerland* (Shuruk), and *Suss v. Germany* (Suss), both also Article 8 cases and both cited in Wetjen. However, while in Wetjen and Suss, the judges have decided a violation of Article 8 has not occurred, they have ruled the opposite in Shuruk.

Let us briefly consider the facts of these cases. In Shuruk, religious parents fight for an extradition of a child. The mother of the child argues that it would be an interference with her right to a family life if the child was to be extradited to the husband, who has joined an ultra-orthodox Jewish movement. A component of the case is an allegation of domestic violence the husband was supposed to perpetrate against his wife. In Wetjen, religious parents used caning (among other methods) as a punishment for their children. The German State has intervened and placed the children in foster care. The parents claim interference with their right to family life. In Suss, German State has denied a divorced father access to his daughter due to the frequent quarrels between the parents during the visits. The father alleges breach of Article 8.

On the surface, the facts are alike, especially




Figure 3: Article 8 and 13 results for Simple, MTL, Joint and Precedent models.

between Shuruk and Wetjen – both cases contain elements of abuse, religion and state intervention. However, to a human lawyer, the distinction between the cases is fairly trivial – the facts are more alike between Wetjen and Suss. In Wetjen, the State is allowed to intervene to protect a child from an abusive ultra-religious parents, which is very similar situation to Suss, where a State is allowed to intervene to protect the child from quarrels between divorced parents.

All our models are exposed to both Shuruk and Suss in the training set. However, for the positive outcome baseline, the information about Suss being related to Article 8 is lost. Conversely, for the negative outcome baseline, the information about Shuruk being related to Article 8 is lost. It is therefore not surprising that the best performing negative⁸ and positive⁹ outcome baseline models both get the Wetjen outcome prediction wrong. On the other hand, the best claim–outcome model, which is trained to learn that both Shuruk and Suss are related to Article 8 via the claim prediction objective, gets Wetjen outcome prediction right.

In conclusion, while a better way of modelling negative outcome, our claim-outcome model is not actually superior in learning anything about the law itself. It simply leverages the fact that outcomes and claims are easy to predict and enforces the relationship between them. To identify the negative outcome with high F_1 will require deeper understanding of law than our models are currently capable of.

9 Related Work

Juris-informatics can trace its origins all the way to the late 1950s (Kort, 1957; Nagel, 1963). The pioneers have used rule based systems to successfully capture aspects of legal reasoning in thou-

⁸Simple Baseline Longformer.

⁹Simple Baseline LEGAL-BERT.

sands of hand crafted rules (Ashley, 1988). Yet these systems have been too brittle to be employed in practice due to the ever-changing rules of law. Especially in common law countries, the majority of law is contained in case law, where cases are transcripts of the judicial decisions, making the law change constantly with every new decision. With the advances of natural language processing (NLP) in the past two decades, the interest in developing NLP applications for the legal domain have been rejuvenated by the research aiming towards more robust models of law. Areas explored include question answering (Monroy et al., 2009), legal entity recognition (Cardellino et al., 2017), text summarisation (Hachey and Grover, 2006), judgement prediction (Xu et al., 2020) and majority opinion prediction (Valvoda et al., 2018).

Our work is similar to the recent study of Chinese law judgement prediction by Zhong et al. (2018) and Xu et al. (2020), who break down the court judgement into the applicable law, charges and terms of penalty. Operating in the civil law system (Germany, China, France etc.), they argue that predicting applicable law is one of the fundamental subtasks, which will guide the prediction for other subtasks. In the context of ECHR law, we argue that legal claims are one such guiding element for outcome prediction. While similar, applicable law and claims are different. In the work above, the judge selects the applicable law from the facts as part of reaching the outcome. This is not the case for ECHR law, or any other precedential legal system known to the authors, where the breach of law is claimed by a lawyer, not a judge.

Another reason that legal claims are not part of the research above is that claims are not available in the Chinese Legal Judgment Prediction task or the large Chinese AI and Law challenge dataset (Xiao et al., 2018). Yet that does not make claims unique for the ECHR law. They are as much part of the ECtHR system as any other precedential legal system. Good candidates for future exploration of the interaction of precedent and claims include the Caselaw Access Project and US Supreme Court caselaw.¹⁰ While ECtHR cases do conveniently provide claim information as part of the judgement, in most legal systems, the information will be submitted directly to the court and to the other party as part of a legal bundle.

¹⁰See US Supreme Court corpus and Caselaw Access Project.

ECtHR dataset has been collected by Chalkidis et al. (2019), who have predicted outcomes of the ECHR law and the corresponding Articles using neural architectures. Our work builds on their research by reinstantiating the outcome prediction task on this dataset to include negative precedent.

10 Conclusion

While positive and negative outcomes are equally important from the legal perspective, the current legal AI research has neglected the latter. Our findings suggest that, at least for current deep learning models, negative outcome is much harder to predict than positive outcome. This has severe implications on how well can the current legal models learn law from legal text. The same models that predict positive precedent with 75.06 F_1 fall short of a random baseline of 10.09 F_1 on the negative outcome prediction task.

We have formulated a theory of why negative outcome prediction is so much harder to learn. Specifically, we suspect that negative outcomes are mostly caused by a judge distinguishing the case from its precedent. This lead us to believe that learning to predict negative outcomes requires more legal understanding than the current models are capable of, which makes it is a particularly attractive task for evaluating progress in legal AI.

While our work significantly improves over the existing models by inducing the relationship between the judge and the lawyer in our claim–outcome model architecture, the best negative outcome prediction model achieves only a third of the performance of the positive outcome one. In future work we hope to study the precedent more closely to build models capable of narrowing this performance gap. One possible avenue would be to relax our Assumption 1 to study the potential relationships between the individual legal Articles.

11 Ethical Considerations

Legal models similar to the ones we study above have been deployed in the real world. Known applications include risk assessment of detainees in the US (Kehl and Kessler, 2017) and sentencing of criminals in China (Xiao et al., 2018). We believe these are not sensible usecases for legal AI. One must not be tempted to think of outcome prediction as equivalent to some medical task, such as cancer detection, with a breach of law seen as a ‘tumour’ that is either there or not. This

is a naive viewpoint that ignores the fact that the legal system is a human construct in which the decision makers, the judges, play a role that can shift the truth, something that is impossible when it comes to natural laws. Put bluntly, if a doctor makes the wrong diagnosis, the patient dies, but if the judge makes the wrong decision, this decision still becomes law.

Herein lies the ethical dubiousness of any attempt at modelling judges using AI. Unlike in the domain of medicine, where identifying the underlying truth is essential for treatment, and thus a successful machine diagnostician is in theory a competition for the human one, in the domain of law the validity of the decision is poised solely on the best intentions of the judge. For some judges this pursuit of the 'right' outcome can go as far as defiance of legal precedent; Lord Denning, whom Margaret Thatcher famously called "*probably the greatest English judge of modern times*" (Stephens, 2009) is one such example. We therefore argue a judge should not be replaced by a machine and caution against the use of our, or any other legal AI model currently available, towards automating the judicial system.

A Note on the Baselines

Comparing the MTL baseline and the joint model, one might come to the conclusion that there is no substantial difference between the models when it comes to predicting positive outcomes. While the joint model outperforms the MTL baseline on eleven out of the twelve experiments we test our models on, the performance on the positive outcome prediction over the outcome corpus is very similar between the two models. None-the-less, there is an important difference between them. The MTL baseline, much like the simple baseline, can predict positive and negative outcome simultaneously for the same Article. This means that in our evaluation, the baseline models can cheat by doubling their predictions. This is another reason that the outcome prediction task needs to consider the legal relationship between the positive and negative outcomes. The lack of this consideration renders both of our baselines fundamentally ill-suited for the task and hence serve only useful for a comparison in our study.

B Glossary & Dataset Examples:

Legal Terms

<i>Claim</i>	The allegation of a breach of law usually put forth by their legal counsel on behalf of the claimant.
<i>Positive Outcome</i>	Claims are assessed by judges in courts of law. If they think a claim is valid, they rule it as successful. The outcome of the case is a victory for the claimant; which we call the positive outcome in this paper.
<i>Negative Outcome</i>	On the other hand, the claimant can be unsuccessful in the court. The judge has decided against them in the court, in favour of the defendant, and we call this the negative outcome in this paper.
<i>Facts</i>	The description of what happened to the claimant. This includes more general descriptions of who they are, circumstances of the perceived violation of their rights and the proceedings in domestic courts before their appeal to ECtHR.
<i>Precedent</i>	Cases that have been cited by the judges as part of their arguments.
<i>Binding</i>	Judges are expected to adhere to the binding rules of law and decide future access accordingly.
<i>Stare Decisis</i>	New cases with the same facts to the already decided case should lead to the same outcome. This is the doctrine of precedent by which judges can create law.
<i>Caselaw</i>	Transcripts of the court proceedings.
<i>ECHR</i>	European Convention of Human Rights, comprises of the Convention and the Protocols to the convention. The Protocols are the additions and amendments to the Convention introduced after the signing of the original Convention.
<i>ECtHR</i>	European Court of Human Rights, adjudicates ECHR cases.
<i>Apply</i>	A judge applies the precedent when she decides on the outcome of a case via an analogy to an already existing case.
<i>Distinguish</i>	Conversely, a judge distinguishes the case from the already existing cases when she believes they are not analogous.

ECtHR Example

Facts	Claims	Positive Outcomes	Negative Outcomes
“Ms Ivana Dvořáčková was born in 1981 with Down Syndrome (trisomy 21) and a damaged heart and lungs. She was in the care of a specialised health institution in Bratislava. In 1986 she was examined in the Centre of Paediatric Cardiology in Prague-Motole where it was established that, due to post-natal pathological developments, her heart chamber defect could no longer be remedied...” for more see Case of Dvoracek and Dvorackova v. Slovakia	Articles: 2, 6, 8, 14	Articles: 2, 6	Articles: 8, 14

C Corpora Statistics:

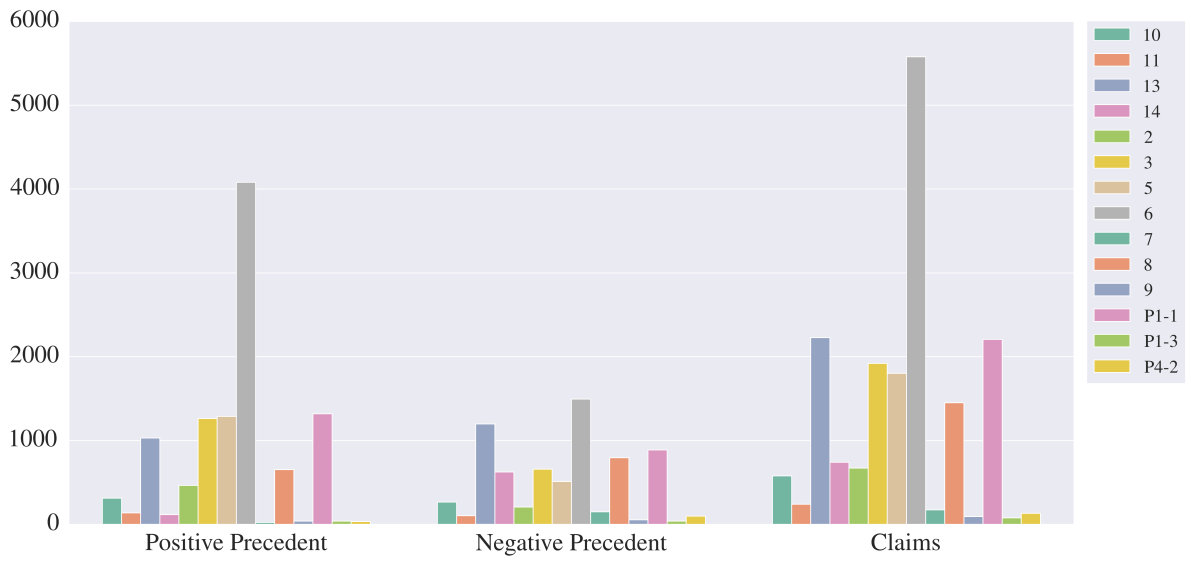


Figure 4: Distribution of Articles over training data in our outcome corpus.




Figure 5: Distribution of Articles over training data in [Chalkidis et al.](#) corpus.

References

- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoŕiuc-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.
- Kevin D. Ashley. 1988. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. Ph.D. thesis, USA. Order No: GAX88-13198.
- Henry Black. 2019. *Black’s Law Dictionary*, 11th edition. Thomson Reuters.
- Cristian Cardellino, Milagro Teruel, Laura Alonso Alemany, and Serena Villata. 2017. Legal NERC with ontologies, Wikipedia and curriculum learning. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 254–259, Valencia, Spain. Association for Computational Linguistics.
- 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. 2021a. 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 Martin Katz, and Nikolaos Aletras. 2021b. LexGLUE: A benchmark dataset for legal language understanding in English. *arXiv*.
- Benjamin Clavié and Marc Alphonso. 2021. The unreasonable effectiveness of the baseline: Discussing SVMs in legal text classification. *arXiv*.
- 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. *The Nature and Authority of Precedent*, chapter 4. Cambridge University Press.
- Kanstantsin Dzehtsiarou. 2017. What is law for the European court of human rights. *Georgetown Journal of International Law*, 49:89.
- ECtHR. 2020. Overview of the case-law of the ECHR. Annual Report 2020 of the European Court of Human Rights, Council of Europe.
- Ben Hachey and Claire Grover. 2006. Extractive summarisation of legal texts. *Artificial Intelligence and Law*, 14(4):305–345.
- D. Kehl and Samuel Ari Kessler. 2017. Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. Esponsive Communities Initiative, Berkman Klein Center for Internet Society, Harvard Law School.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Fred Kort. 1957. Predicting supreme court decisions mathematically: A quantitative analysis of the “right to counsel” cases. *American Political Science Review*, 51(1):1–12.

- 1400 Grant Lamond. 2016. [Precedent and Analogy](#)
1401 [in Legal Reasoning](#). In Edward N. Zalta,
1402 editor, *The Stanford Encyclopedia of Philos-*
1403 *ophy*, Spring 2016 edition. Metaphysics Re-
1404 search Lab, Stanford University. 1450
- 1405
1406 Reed C. Lawlor. 1963. [What computers can do:](#)
1407 [Analysis and prediction of judicial decisions](#).
1408 *American Bar Association Journal*, 49(4):337–
1409 344. 1451
- 1410 Yonatan Lupu and Erik Voeten. 2010. [The role](#)
1411 [of precedent at the European court of human](#)
1412 [rights: A network analysis of case citations](#).
1413 *OpenSIUC*. 1452
- 1414
1415 Alfredo Monroy, Hiram Calvo, and Alexander
1416 Gelbukh. 2009. [NLP for shallow question](#)
1417 [answering of legal documents using graphs](#). vol-
1418 ume 5449, pages 498–508. 1453
- 1419
1420 Stuart S. Nagel. 1963. [Applying correlation anal-](#)
1421 [ysis to case prediction](#). *Texas Law Review*,
1422 42:1006. 1454
- 1423
1424 Adam Paszke, Sam Gross, Francisco Massa,
1425 Adam Lerer, James Bradbury, Gregory
1426 Chanan, Trevor Killeen, Zeming Lin, Natalia
1427 Gimelshein, Luca Antiga, et al. 2019. [Py-](#)
1428 [torch: An imperative style, high-performance](#)
1429 [deep learning library](#). In *Advances in Neu-*
1430 *ral Information Processing Systems*, pages
1431 8026–8037. 1455
- 1432
1433 Charles Stephens. 2009. *The Jurisprudence of*
1434 *Lord Denning: A Study in Legal History, in*
1435 *Three Volumes*. Cambridge Scholars Publish-
1436 ing, Cambridge. 1456
- 1437
1438 Josef Valvoda, Tiago Pimentel, Niklas Stoehr,
1439 Ryan Cotterell, and Simone Teufel. 2021. [What](#)
1440 [about the precedent: An information-theoretic](#)
1441 [analysis of common law](#). In *Proceedings of the*
1442 *2021 Conference of the North American Chap-*
1443 *ter of the Association for Computational Lin-*
1444 *guistics: Human Language Technologies*, pages
1445 2275–2288, Online. Association for Computa-
1446 tional Linguistics. 1457
- 1447
1448 Josef Valvoda, Oliver Ray, and Ken Satoh. 2018.
1449 [Using agreement statements to identify major-](#)
1450 [ity opinion in UKHL case law](#). In *Legal Knowl-*
1451 *edge and Information Systems*, pages 141–150.
1452 IOS Press. 1458
- 1453
1454 Ashish Vaswani, Noam Shazeer, Niki Parmar,
1455 Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
1456 Łukasz Kaiser, and Illia Polosukhin. 2017. [At-](#)
1457 [tention is all you need](#). In *Advances in Neural*
1458 *Information Processing Systems*, pages 5998–
1459 6008. 1454
- 1460
1461 Thomas Wolf, Lysandre Debut, Victor Sanh,
1462 Julien Chaumond, Clement Delangue, Anthony
1463 Moi, Pierric Cistac, Tim Rault, Remi Louf,
1464 Morgan Funtowicz, Joe Davison, Sam Shleifer,
1465 Patrick von Platen, Clara Ma, Yacine Jernite,
1466 Julien Plu, Canwen Xu, Teven Le Scao, Sylvain
1467 Gugger, Mariama Drame, Quentin Lhoest, and
1468 Alexander Rush. 2020. [Transformers: State-](#)
1469 [of-the-art natural language processing](#). In *Pro-*
1470 *ceedings of the 2020 Conference on Empirical*
1471 *Methods in Natural Language Processing: Sys-*
1472 *tem Demonstrations*, pages 38–45, Online. As-
1473 sociation for Computational Linguistics. 1456
- 1474
1475 Chaojun Xiao, Haoxi Zhong, Z. Guo, Cunchao
1476 Tu, Zhiyuan Liu, M. Sun, Yansong Feng, Xi-
1477 anpei Han, Z. Hu, Heng Wang, and J. Xu. 2018.
1478 [CAIL2018: A large-scale legal dataset for judg-](#)
1479 [ment prediction](#). *ArXiv*, abs/1807.02478. 1460
- 1480
1481 Nuo Xu, Pinghui Wang, Long Chen, Li Pan, Xi-
1482 aoyan Wang, and Junzhou Zhao. 2020. [Distin-](#)
1483 [guish confusing law articles for legal judgment](#)
1484 [prediction](#). *arXiv preprint arXiv:2004.02557*. 1461
- 1485
1486 Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chao-
1487 jun Xiao, Zhiyuan Liu, and Maosong Sun.
1488 2018. [Legal judgment prediction via topolog-](#)
1489 [ical learning](#). In *Proceedings of the 2018 Con-*
1490 *ference on Empirical Methods in Natural Lan-*
1491 *guage Processing*, pages 3540–3549, Brussels,
1492 Belgium. Association for Computational Lin-
1493 guistics. 1462
- 1494
1495 Bostjan Zupancic. 2016. [In the context of the](#)
1496 [common law: The European court of human](#)
1497 [rights in Strasbourg transcript](#). Transcript of a
1498 lecture given at Gresham College. 1463