The Research Frontier: AI and Legal Reasoning

Alina Petrova, John Armour, Thomas Lukasiewicz, YordanYordanov

AI for English Law Webinar: LawTech – Progress and Challenges

15 September 2021

INDUSTRIAL STRATEGY | UK Research and Innovation

UNIVERSITY OF OXFORD

Motivation

- Automated solutions that scale can **dramatically** lower user costs

  Example: First mark-up of a standard type of contractual agreement, which might be charged at £1,000+ by a lawyer working in a top-tier law firm, can today be done by an AI system for less than £1.

⇒For citizens, facilitate access to justice

⇒For businesses, lower costs

Research Question

- How can we harness AI to give (better) guidance on dispute resolution?

Potential use-cases

Advice

What are my legal rights?

What should I do?

Strategy

Should I settle/fund this suit?

For how much?

Resolution

Which party should succeed?

For what reasons?

Review

Should a first-instance decision be subject to appeal or review?

State of the art

Simple outcome prediction

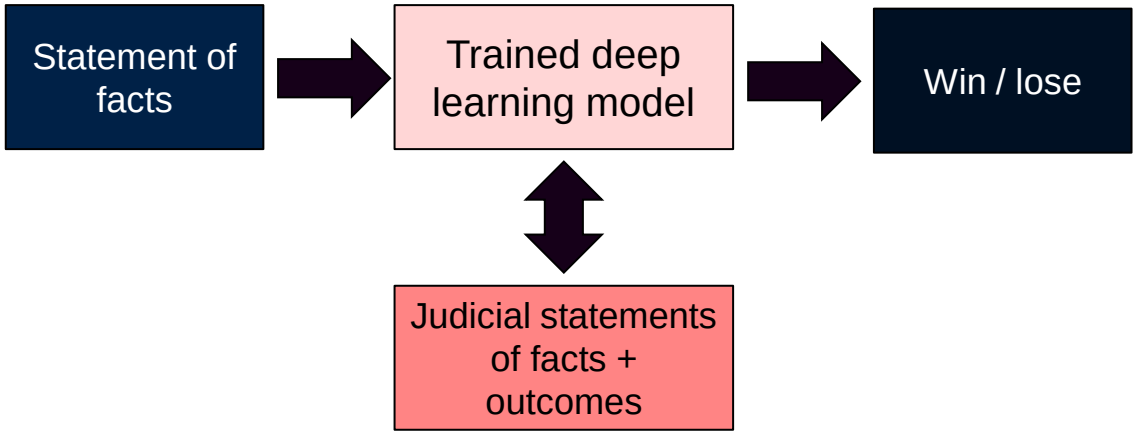

- Prediction = outcome: claimant win/lose (eg Aletras et al, 2016)
- Around 70-80% accuracy can be achieved if exclude cases with multiple issues.
- Applied to ECtHR caselaw (published in structured format)
- No explanation for predicted outcome

⇒ **Research frontier**: how to apply to caselaw not published in a structured format? How to explain predictions?

INDUSTRIAL STRATEGY | UK Research and Innovation

Cause of action identification

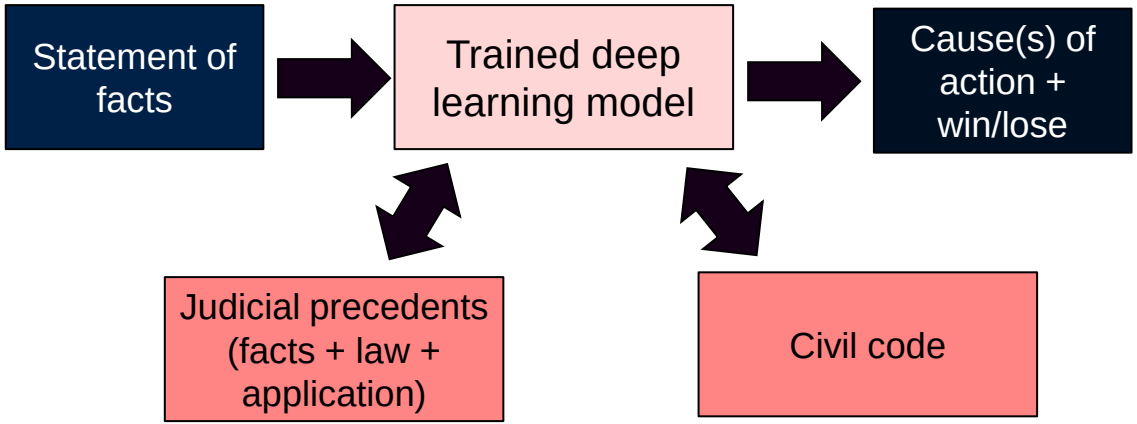

- Prediction: which causes of action are relevant
- So far, only implemented with Chinese legal data (using civil code) (Yang et al, 2019)
- Civil code acts as authoritative “map” of legal causes of action

⇒ **Research frontier**: Implement in a common law system.

Potential use-cases

Advice

What are my legal rights?

What should I do?

Strategy

Should I settle/fund this suit?

For how much?

Resolution

Which party should succeed?

For what reasons?

Review

Should a first-instance decision be subject to appeal or review?

Data

Approaches to digital access to caselaw

China:

Over 88m court documents freely available online at central repository; state sponsorship of legal AI competitions

(CAIL2018, 2019)

Data

USA: many court pre
available online;
centralised reposit
PACER has dock
information for many
but charges $0.10/p

Harvard Law Scho
Caselaw Access Proje
6.4m US cases

n on using
data in
tics (Loi
t 33)),
o 5 years
nt

The Law Society Gazette

Contact us

NEWS | ANALYSIS | LAW | PRACTICE | IN-HOUSE | PEOPL

NEWS

BAILII grants access to judgments for mass AI analysis

By Michael Cross | 14 December 2020

4 Comments

England and Wales court judgments are for the first time being opened to mass analysis by artificial intelligence, the Gazette can reveal. Under an agreement announced today, the British and Irish Legal Information Institute (BAILII) has granted academics at Oxford University bulk access to its database of 400,000 judgments for research purposes.

Until now, BAILII, the principal online repository of case law, has refused to allow its database to be downloaded in bulk or 'scraped' by software. The fear is that this would enable the development of AI systems for predicting the outcome of cases on a judge-by-judge basis. It is understood that the agreement between BAILII and Oxford's AI for Law project excludes such use. Rather, the agreement will 'unlock new research insights into English case law' and help the development of ways to improve access to legal information, a statement said. The negotiations with BAILII took over a year and encompassed review by stakeholders at the Ministry of Justice and the judiciary, it added.

Most open

Least open

INDUSTRIAL STRATEGY | UK Research and Innovation

Data Challenges

Hope:

Case description: On July 7, 2017, when the defendant Cui XX was drinking in a bar, he came into conflict with Zhang XX…… After arriving at the police station, he refused to cooperate with the policeman and bited on the arm of the policeman……

Result of judgment: Cui XX was sentenced to *12* months imprisonment for *creating disturbances* and *12* months imprisonment for *obstructing public affairs*……

- Charge#1 creating disturbances term 12 months
- Charge#2 obstructing public affairs term 12 months

Chen et al. *Charge-Based Prison Term Prediction with Deep Gating Network*. 2019

Reality:

INDUSTRIAL STRATEGY | UK Research and Innovation

Challenges (cont'd)

Data challenges:

- No access to claim forms/pleadings or settled cases.
- Lack of metadata and annotations.
- Information extraction: outcomes, facts, legal reasons.

Modelling challenges:

- Manual annotation leads to small sample size.
- Most data is out-of-sample: different judges, courts, countries.
- Lack of first instance cases; appellate data means less informative.

INDUSTRIAL STRATEGY | UK Research and Innovation

Experiments and Results

Predicting Outcome + Reasons

Step 1: Annotation

- *facts* (what initially happened)
- *procedural history* (inc. what lower court ruled)
- *relevant precedents* (which prior precedents applied)
- *application of law* (how law is applied to the facts)
- *outcome*
- *other discussion*

INDUSTRIAL STRATEGY | UK Research and Innovation

Manual annotation

- Team of 3 RAs (law research students)
- Multiple iterations to converge on equivalent coding
- Coding protocols continuously developed
- Weekly meetings / problem workshops

- Working total: c. 500 cases

Contributions

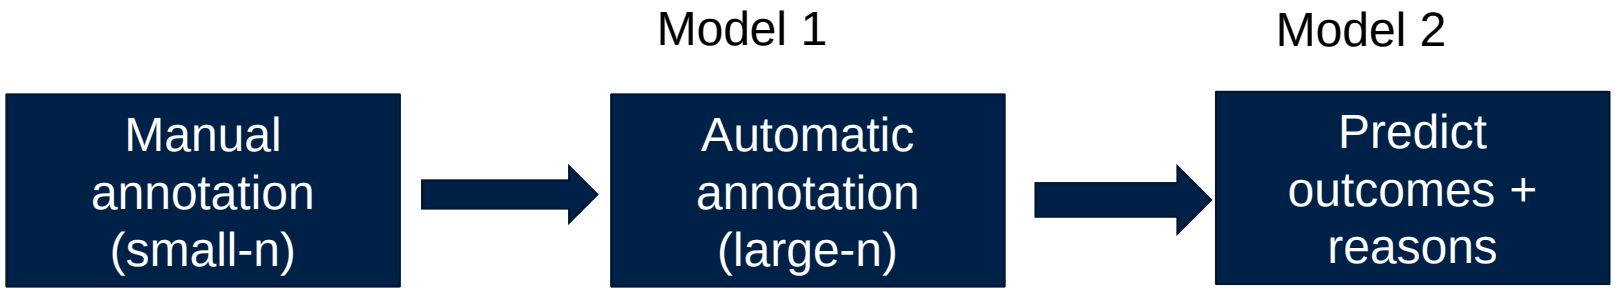

Automating annotation: outcomes

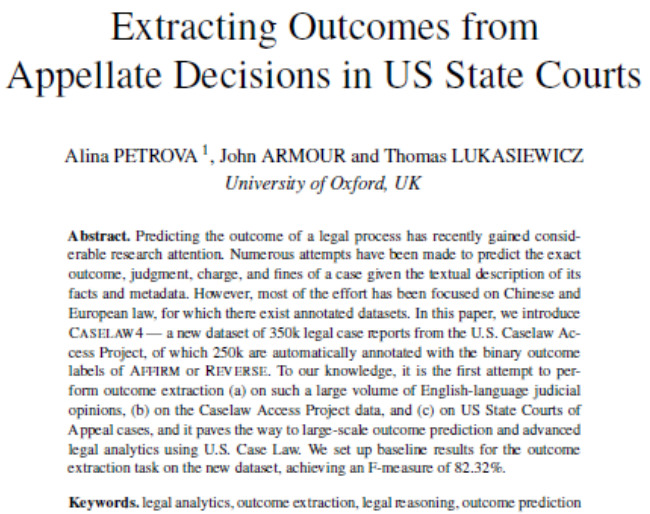

Extracting Outcomes from Appellate Decisions in US State Courts

Alina PETROVA [1], John ARMOUR and Thomas LUKASIEWICZ
University of Oxford, UK

Abstract. Predicting the outcome of a legal process has recently gained considerable research attention. Numerous attempts have been made to predict the exact outcome, judgment, charge, and fines of a case given the textual description of its facts and metadata. However, most of the effort has been focused on Chinese and European law, for which there exist annotated datasets. In this paper, we introduce CASELAW4 — a new dataset of 350k legal case reports from the U.S. Caselaw Access Project, of which 250k are automatically annotated with the binary outcome labels of AFFIRM or REVERSE. To our knowledge, it is the first attempt to perform outcome extraction (a) on such a large volume of English-language judicial opinions, (b) on the Caselaw Access Project data, and (c) on US State Courts of Appeal cases, and it paves the way to large-scale outcome prediction and advanced legal analytics using U.S. Case Law. We set up baseline results for the outcome extraction task on the new dataset, achieving an F-measure of 82.32%.

Keywords. legal analytics, outcome extraction, legal reasoning, outcome prediction

Petrova et al (2020)

- 500 randomly-selected cases manually annotated for outcome (affirm, reverse, mixed)
- Deep learning model trained using hand-coded sample
- 250,000 cases automatically annotated with 82.3% accuracy

Automating annotation: facts

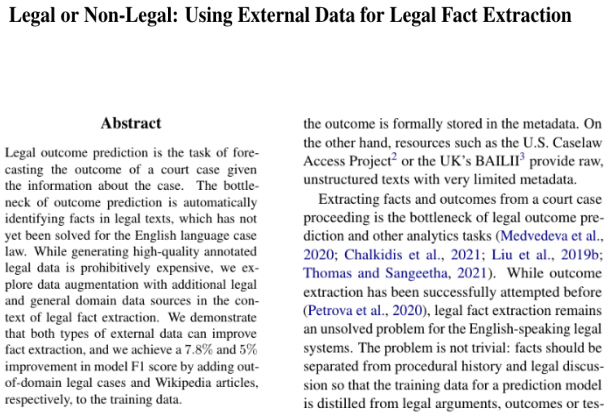

Legal or Non-Legal: Using External Data for Legal Fact Extraction

Abstract

Legal outcome prediction is the task of forecasting the outcome of a court case given the information about the case. The bottleneck of outcome prediction is automatically identifying facts in legal texts, which has not yet been solved for the English language case law. While generating high-quality annotated legal data is prohibitively expensive, we explore data augmentation with additional legal and general domain data sources in the context of legal fact extraction. We demonstrate that both types of external data can improve fact extraction, and we achieve a 7.8% and 5% improvement in model F1 score by adding out-of-domain legal cases and Wikipedia articles, respectively, to the training data.

the outcome is formally stored in the metadata. On the other hand, resources such as the U.S. Caselaw Access Project[2] or the UK's BAILII[3] provide raw, unstructured texts with very limited metadata.

Extracting facts and outcomes from a court case proceeding is the bottleneck of legal outcome prediction and other analytics tasks (Medvedeva et al., 2020; Chalkidis et al., 2021; Liu et al., 2019b; Thomas and Sangeetha, 2021). While outcome extraction has been successfully attempted before (Petrova et al., 2020), legal fact extraction remains an unsolved problem for the English-speaking legal systems. The problem is not trivial: facts should be separated from procedural history and legal discussion so that the training data for a prediction model is distilled from legal arguments, outcomes or tes-

Petrova et al (2021)

- 70 negligence cases from CAP, annotated for legal facts
- Data augmentation with ECHR cases, news reports, Wikipedia articles
- Fact extraction accuracy improves by 7.8% with additional legal data and by 5% with non-legal data

INDUSTRIAL STRATEGY | UK Research and Innovation

Predicting Outcome + Reasons

Step 2: Giving context via:

- Legal issues
  what legal issues are relevant to the facts? (Restatement)

- Prior caselaw
  which prior cases are relevant to the facts?

- Reasons: application of law to facts
  what reasons are given for application of law to facts in similar cases?

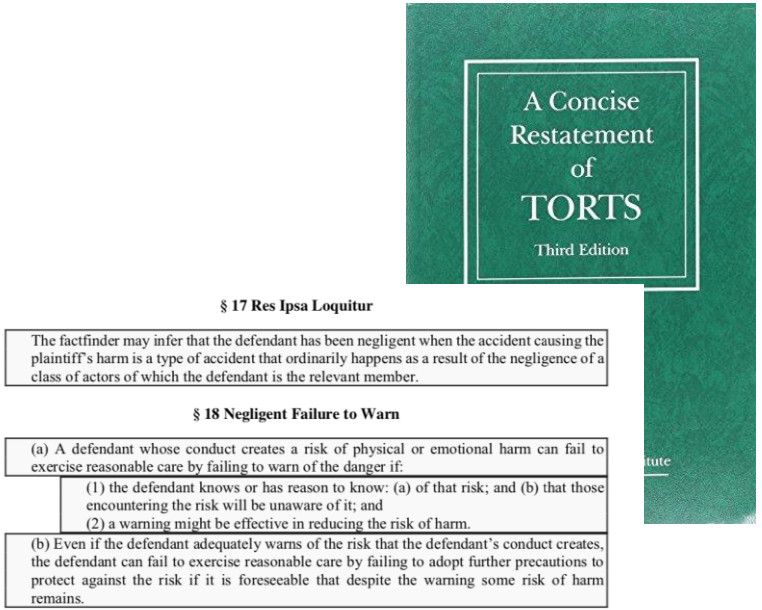

INDUSTRIAL STRATEGY | UK Research and Innovation

Predicting Outcome + Reasons

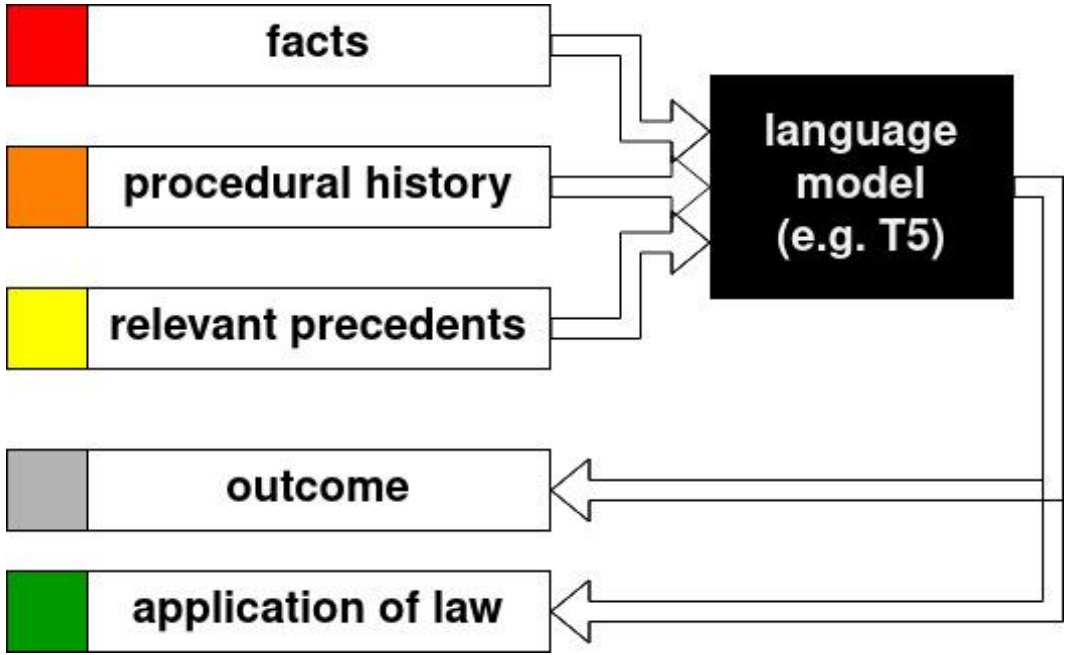

The language model is pre-trained on a large corpus of multi-domain text.

Relevant precedent retrieval could be an intermediary step.

Applications?

Potential use-cases