Finding Case Law: Leveraging Machine Learning Research to Enhance Public Access to UK Judgments

Amy Conroy[1,2,§], Editha Nemsic[1,3,§], Daniel Hoadley[1] and Imane Hafnaoui[1,4]

[1]Mishcon de Reya LLP, London, WC2B 6AH, United Kingdom
[2]University of Bristol, Bristol, BS8 1TH, United Kingdom
[3]University College London, London, WC1E 6BT, United Kingdom
[4]Queen Mary University of London, London, E1 4NS, United Kingdom

Abstract

Once ranked last in Europe for public access to judgment data, the United Kingdom has taken large strides in recent years to improve the accessibility of judgments. This paper discusses how the new platform from The National Archives, *Find Case Law*, was developed for the publication of UK judgments; in particular how we created the engine responsible for the enrichment of judgment text. We argue that the new system is necessary to address existing issues with the accessibility of judgment data, and if the platform were to leverage the abundance of research conducted in areas such as legal text classification, summarisation, and entity recognition, the UK could quickly become a world leader for public accessibility of judgments. We develop a proof of concept system, *MyJudgments*, that demonstrates a potential direction for development. Whilst it is early days, the launch of Find Case Law provides a unique opportunity to remind ourselves of the opportunities machine learning presents for broadening the accessibility of judgments to new users and expanding their utility for novel use-cases. To do this, we review existing research performed on UK judgment data and suggest how the various strands could practically be integrated into a case law publication system.

Keywords

Case law, machine learning, judgment publication

1. Introduction

The United Kingdom has long fallen behind its European counterparts in regards to public access to judgment data, made clear by a 2018 report from the European Commission which ranked the UK last in comparison to other European countries [1]. Judgment dissemination in the UK has traditionally been carried out by commercial publishers who selectively published precedent-setting case law in the form of law reports on a paid-subscription basis. Since the early 2000s and until relatively recently, the de facto official online source of publicly accessible

[§]These authors contributed equally to this work.

Joint Proceedings of ISWC2022 Workshops: the International Workshop on Artificial Intelligence Technologies for Legal Documents (AI4LEGAL) and the International Workshop on Knowledge Graph Summarization (KGSum) (2022)

amy.conroy@mishcon.com (A. Conroy); editha.nemsic@mishcon.com (E. Nemsic); daniel.hoadley@mishcon.com (D. Hoadley); imane.hafnaoui@mishcon.com (I. Hafnaoui)
https://www.amyconroy.co.uk (A. Conroy)
0000-0002-4030-0337 (A. Conroy)

judgments in the UK was BAILII[1]. BAILII's case law coverage is not comprehensive. Recent work comparing BAILII's coverage of judicial review judgments with that provided by the commercial research platform vLex Justis identified a significant gap in public access throughout the period measured [2]. The authors of that study, building on earlier analyses [3, 4] conjectured that the gap in public access is attributable to the complicated court recording and transcription regime under which privately-owned transcription agencies convert oral judgments into written form for a fee. In contrast with its commercial counterparts, BAILII lacks the funds to obtain these transcripts at commercial rates, thereby rendering a substantial portion of judgments accessible only behind a paywall. This state of affairs is particularly problematic in common law jurisdictions like the UK where judgments constitute a primary source of law. To address this, in 2021 the Ministry of Justice announced that cases 'of legal significance' would be published on a new platform called Find Case Law, built and maintained by The National Archives (TNA) [5].

This paper discusses the development of the government-backed *Find Case Law* XML enrichment engine, including the introduction of open source annotation pipelines for case law citations, abbreviations, and legislative instruments. We provide an overview of where public access to judgments currently stands in the UK, and where it still has the potential to go if existing machine learning research was incorporated into a case law publication platform. We review existing research into the use of computational methods on UK judgments, such as rhetorical role labelling and summarisation. To visualise how this research can be leveraged in a practical way to improve public access to judgment data in the UK, we develop a proof of concept system ('MyJudgments') that incorporates the work delivered by the Find Case Law platform and the other machine learning research[2].

2. Background

2.1. Linked Data and LegalDocML

The majority of research surrounding the development of a Case Law publication system has focused on the optimisation of the data within the judgment itself. This includes the addition of links between different Government applications, such as a legislation and case law database [6]. Adding direct links between different sources of legal principles allows those who are non-experts to interact with the data in the same way that an expert might infer the links between sources simply by reading the judgment text. An attempt to standardise this was introduced by the European Union in the form of the European Law Identifier (ELI) and the European Case Law Identifier (ECLI) [6].

In order to improve the accessibility of judgments in the UK, the Government has chosen to adopt the Legal Document Mark-Up Language (LegalDocML) [7]. LegalDocML is a standard developed for legal documents, including legislative, parliamentary, and judicial documents. Leveraging LegalDocML allows the different functions of Government to interact by sharing common document and metadata points. The implementation of the LegalDocML standard,

[1] https://www.bailii.org

[2] https://github.com/mdrresearch/MyJudgments

particularly with reference to case law and legislation links, will be discussed further in the next section.

2.2. Machine Learning Research on Judgments

There have been a number of various research strands relating to applying machine learning and other computational techniques to judgment data, often with the motivation of improving public access to the data. This includes research into automatic summarisation systems both in the UK [8, 9] and other jurisdictions [10, 11], the analysis of agreement statements between judges to identify the majority opinion [12], and other text classification experiments [13]. Leveraging this research to enhance existing case law publication systems has the ability to improve the accessibility of judgments, and put experts and non-experts on a level playing field. We explore UK-focused research in further detail in the below sections.

2.3. Motivation

Our motivation for this paper is to review the available machine learning (ML) research on UK judgment data and demonstrate how the current research can be leveraged to improve the understanding of judgments, and thus public access to case law. It is important to focus on research undertaken on UK judgment data given the way that judgments are drafted varies between jurisdictions.

Existing research has largely focused on extracting information from judgments but less so on knowledge presentation techniques for legal documents. The combination of the release of the Find Case Law service in the UK and recent advancements within legal ML research suggest it is the perfect time to focus on combining Government case law platforms with the work done in academia. We propose a proof of concept system that we suggest can easily make use of existing computational techniques to surface valuable judgment information to improve access to judgments for the average citizen. For the purposes of our platform, we suggest that the average citizen is an ordinary, non-legally trained, user without access to the paid-for private case law services typically available for commercial legal teams.

2.4. Considerations

The development of any system for public use must be balanced with the needs and wants of the public as well. A recent report by the Legal Education Foundation aimed to understand how members of the public viewed commercial access to judgments and data in court records [14]. The key findings indicated that 'respondents overwhelmingly found it important' for there to be controls around who can access court data, how they access it, and what they can do with it. Thus, any system developed to aid the public in understanding and accessing judgments needs to be balanced against the wishes of the public. In addition to the above considerations, any work in this area must be balanced against The National Archives' own licensing restrictions. Any computational analysis on judgments from TNA's Find Case Law system can not take place until a transactional licence is obtained[3].

[3] https://caselaw.nationalarchives.gov.uk/transactional-licence-form

3. UK Judgment Enrichment Pipeline

The UK's 'Find Case Law' service went live in April 2022[4], and consisted of a publishing service, a public facing user interface and enriched judgment XML content downloadable from the website. This paper focuses on the development of the XML enrichment engine, rather than the user interface. The enrichment engine consists of five separate annotators; for case law citations, legislation, legislation provisions, oblique references to legislation, and abbreviations. The development of the annotators and resulting output, which uses LegalDocML format, will be explained briefly in the following section, the code for which is available on Github[5].

3.1. Case Law Annotator

The first annotator in the pipeline is the Case Law Annotator, which detects both well-formed and malformed references to UK judgments and links them to the corresponding judgment available on the Find Case Law website. It was important to be able to detect both well-formed and malformed references, as these references are often written incorrectly in UK judgments due to the specificity that is required and the variation of citations between the different courts.

In order to detect malformed citations, a rule-based approach was used where each rule represented a well-formed citation or a sub-set of the most common malformed versions of the citation. These were then stored as rules in spaCy's EntityRuler [15]. The XML that is wrapped around the identified citation includes the canonical, or well-formed, citation and the link to the case. An example of this is: `<ref href="https://caselaw.nationalarchives.gov.uk/ewca/civ/2021/1308"uk:canonical="[2021] EWCA Civ 1308"uk:isneutral="true"uk:type="case"uk:year="2021">2021 EWCA.Civ 1308</ref>`.

3.2. Abbreviation Annotator

The purpose of the abbreviation annotator is to detect abbreviations and resolve the short form (for example, *HRA*) to the long form (*Human Rights Act*). In LegalDocML, this is represented by the following: `<abbr title="Human Rights Act">HRA</abbr>`. We adapted the abbreviation detector from the Blackstone library[6], which itself was an adaptation of scispaCy [16].

The abbreviation detector previously worked by identifying items in brackets and walking backwards to see if the preceding words started with the same letters. In order to account for the way in which traditional abbreviations are defined in UK judgments, we constricted this to apply only where there were brackets and then quotations around the short form. This ensures that only abbreviations of courts or legislation, for example, are detected rather than information in brackets that are not traditional abbreviations (such as an alternative defendant name).

[4] https://caselaw.nationalarchives.gov.uk
[5] https://github.com/nationalarchives/ds-caselaw-data-enrichment-service
[6] https://github.com/ICLRandD/Blackstone

3.3. Legislation Annotators

There are three different annotators that are used to link to relevant legislation as referenced in the judgment. The first legislation annotator in the pipeline applies LegalDocML to the primary legislation referred to. For example, `<ref href="http://www.legislation.gov.uk/id/ukpga/2006/46/"uk:canonical="2006 c. 46"uk:type ="legislation">Companies Act 2006</ref>`, links to the relevant legislation instrument on the legislation.gov.uk website. In order to identify the correct legislation, the annotator uses a combination of exact string and fuzzy matching that references a lookup table of existing Acts. The table is updated every seven days by querying a www.legislation.gov.uk SPARQL endpoint.

In addition to references to the Act itself, we implemented a legislation provision annotator that identifies and links to specific sections of the legislation. An example of this is: `<ref href="http://www.legislation.gov.uk/id/ukpga/2006/46/section/17"uk:canonical="2006 c. 46 s. 17"uk:type="legislation">section 17</ref>`. The final legislation annotator identifies oblique references (such as *this Act* or *the Act*) and links them to the relevant piece of legislation. In LegalDocML, this is represented as `<ref href="http://www.legislation.gov.uk/id/ukpga/1972/68" uk:canonical="1972 c. 68"uk:type="legislation">this Act</ref>`.

The legislation provision and oblique reference annotators were implemented using similar methods. Using the previously enriched judgment, we extracted sentences where a piece of primary legislation had been detected. When we identify reference to an oblique reference or provision using regex, we use the location of the citation to find the closest piece of legislation within a certain threshold. Where sections or oblique references are re-defined to a difference piece of legislation, the subsequent references will be linked to the newly referenced legislation.

4. Application Development and Legal Machine Learning Research

In this section we review existing machine learning research performed on UK judgment data, and suggest how this can be incorporated into a case law publication platform to enhance the public's experience interacting with judgments. In order to do this, we develop a proof of concept system, *MyJudgments* as shown in Figure 1, that provides a simple interface to demonstrate how existing lines of machine learning research could be incorporated into a system that provides value to the end user. With little effort, we demonstrate it is possible to expose levels of detail and additional insight into judgments that are typically only available to those with licenses to private commercial judgment products, or to those who are legally trained and able to infer the contextual information.

Before exploring existing research it's important to explain the proof of concept system, *MyJudgments*. It is a simple user interface designed to be complementary to the Find Case Law platform. It was built with a React back-end, allowing the user to view judgments with access to an additional layer of informative details extracted from the judgment and surfaced in a user-friendly way. The core purpose of the application is to make judgments available in a format that is easy to navigate and digest for the average, non-legally trained, user.

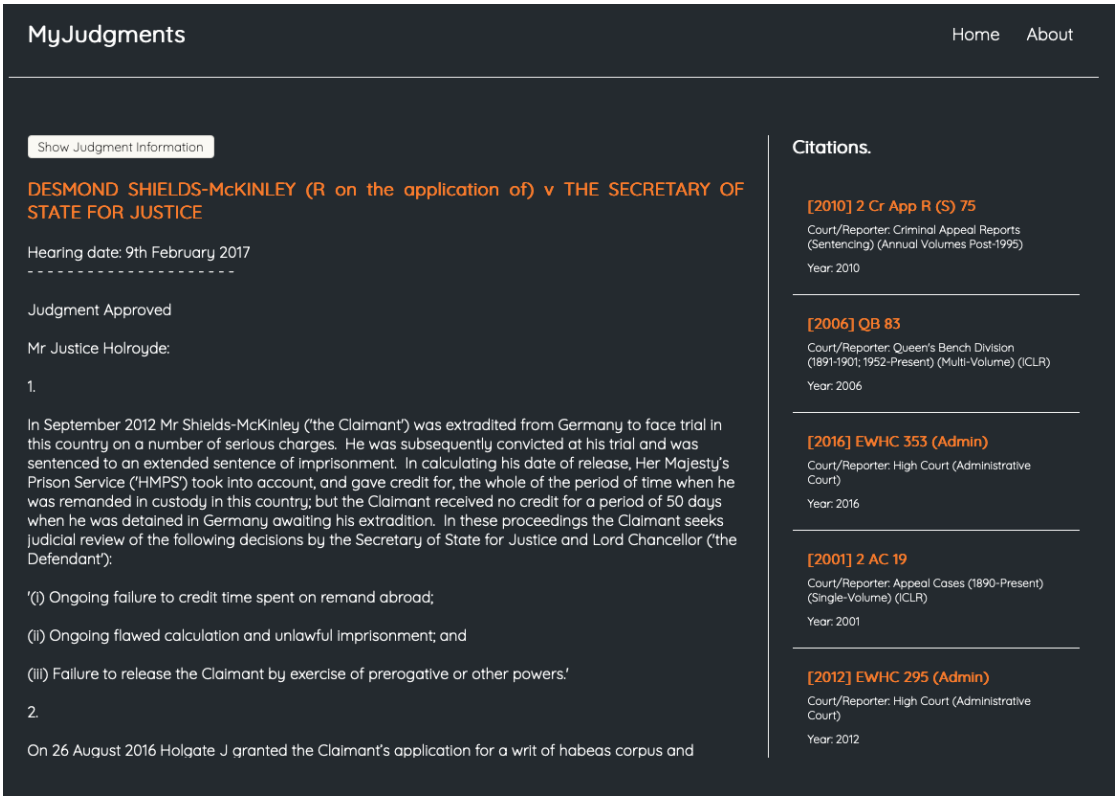

Figure 1: *MyJudgments* view of an example case.

4.1. LegalDocML

In the first instance, we use the enrichment engines provided via the open source Find Case Law Github repository to annotate case law citations, legislation citations and abbreviations. Although the labelling of case law and legislation provisions with LegalDocML result in linked references to the relevant citations, which itself is improving the user experience of interacting with the judgment, we suggest it can be taken further by exposing the underlying information in a visual and interactive way.

We leverage the information provided in the XML to expose the number of case law and legislation citations, a definition key to the abbreviations used within the judgment, and a list of the cases cited within a judgment. Figure 3 displays an example judgment with a citation to a provision within a legislative instrument at the bottom of the page. On the right-hand side there is a hyperlinked list of case law citations that are referenced within the current judgment, allowing them to navigate to the cited case law. For those unfamiliar with certain cases, having instant access to the available citations allows users to understand the precedent that influenced the decision. In a common law system, understanding the case law cited within a given judgment is imperative to comprehend the law itself.

In addition, legal judgments are drafted in a way that makes it difficult for those without

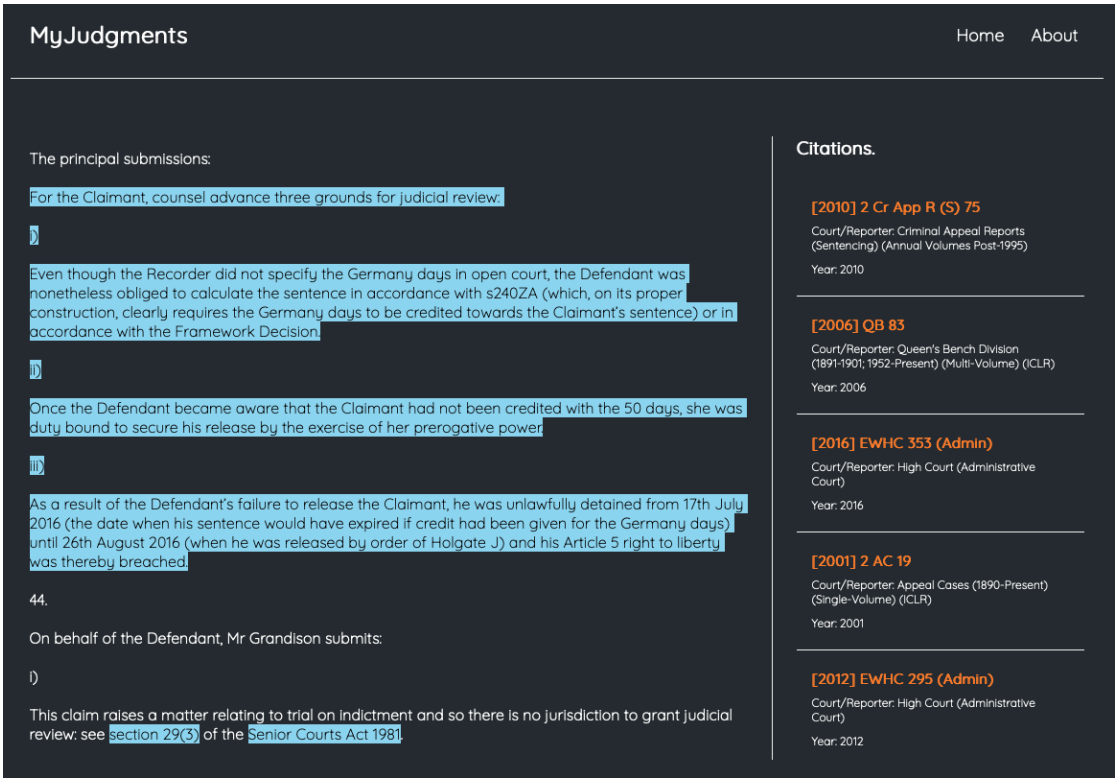

Figure 2: Example judgment with highlighted sentences containing the grounds of the claim and a citation to a legislative instrument.

traditional legal training to quickly grasp the decision and other subtle contextual information, that someone with legal training may easily understand. A common example of this is abbreviations, which are used frequently in judgments, particularly for things such as courts and other legal terms. While those who practice or work in the legal sector are familiar with common abbreviations, others are likely to have to frequently refer back to where the terms were first defined. By providing a definition key that simply extracts the list of abbreviations from the XML of the judgments, as shown in Figure 3, we are exposing readily available information that is hidden within the XML and making it easier for the user to interact with the contents of the judgment.

4.2. Rhetorical Roles

Research into the automatic classification of rhetorical roles on legal judgments across jurisdictions has been plentiful. In the UK, much of the research into rhetorical role classification stems from the early 2000s work by Hachey and Grover [17, 8]. Hachey and Grover's rhetorical role annotator labelled sentences with a label of FACT, PROCEEDINGS, BACKGROUND, FRAMING, DISPOSAL, TEXTUAL, or OTHER. They experiment with different machine learning techniques to assign the roles on a sentence-level. Their best result was a 60.6% F-Score with a support

vector machine classifier.

Bhattacharya et al. used deep learning techniques to classify rhetorical roles on a sentence-level on UK judgments [18]. They found that neural methods such as Hierarchical BiLSTM architectures performed better compared to other ML techniques. They also found that it was better to train models on data from the target jurisdiction, underlying the need for further legal ML research on UK judgments.

Much of the research surrounding the classification of rhetorical roles in judgments has suggested that there would be a large benefit if the roles were to be exposed to the end-user. We demonstrate this in Figure 3, where the highlighted block of sentences has been automatically labelled with the rhetorical role of 'grounds'. To label the sentences we use a light-weight decision tree classifier, trained on manually labelled grounds sentences in judicial review judgments. This allows the user to quickly locate the grounds of the judgment and understand the context of the sentences with respect to the rest of the judgment, in a clear and visual way. The same exercise could be repeated with rhetorical roles such as 'fact' or 'background', for which classifiers have been built within a legal context [9].

4.3. Summarisation

The automatic summarisation of judgments would allow a user to have immediate access to summaries of newly released judgments as well as judgments from lower courts, which typically do not have dedicated individuals to write manual summaries. Using machine learning techniques to generate these summarises in combination with a judgment publication system could also allow for the user to customise the summaries to their desired length, or with the contextual content they require.

The research undertaken by Hachey and Grover mentioned above fed into their automatic summarisation system, the SUM system [8]. They used the rhetorical role classifier in addition to a relevance classifier, which classified sentences as relevant or not. They selected the most relevant sentences from across the respective rhetorical roles to automatically create a summary reflecting the leading manually written UK judgment summaries. Ray et al. built on their research with the SUMO system, implementing a conditional random fields classifier to perform the rhetorical and relevance classifications [9].

In Figure 3, we show how a summary can be included when a user selects 'Show Judgment Information'. Currently a limited number of judgments in the UK have publicly available manually-written summaries. By integrating an automatic case summarisation system we allow for users to quickly gain an understanding of key issues and outcome of all relevant cases instantaneously, rather than being limited to cases deemed legally significant or to cases from higher courts which have manually-written summaries. This eases the understanding gap that might otherwise exist, as well as the access gap to those who have access to paid-for subscriptions to obtain case summaries.

5. Future Work

The proof of concept, *MyJudgments*, presented in this paper was built in a short amount of time with limited resources. It is intended to demonstrate the opportunities existing research has

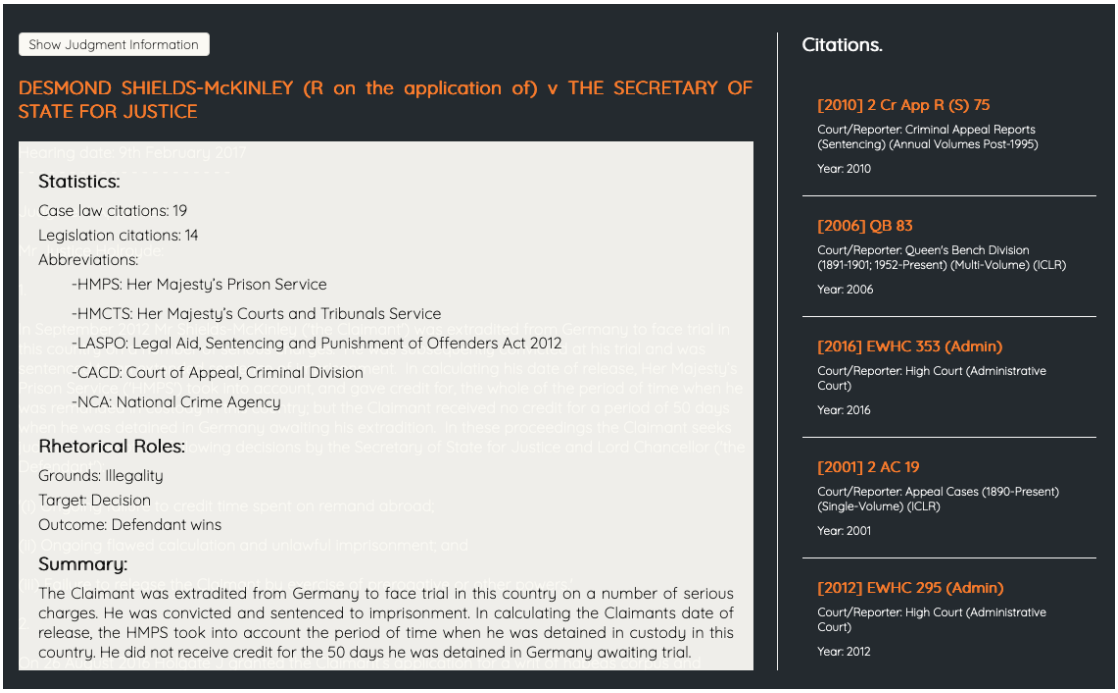

Figure 3: Example judgment with an expanding overlay containing an overview of critical information about the judgment including citation statistics, abbreviations, rhetorical role categories and a summary.

opened up for better access to public judgements. While our system is a solid starting point we recognise a number of potential improvements.

In the first instance, we would like to engage a group of users that includes those that have no legal training through to practising solicitors to understand the value that they would gain from the suggested features. The immediate barrier is the effort and resources required to undertake a software development project of this scale. However, the immediate aim should be a collaboration between researchers and those working on open access case law publication systems.

Examples include a search functionality across the whole system that allows searching for judgment titles, neutral citations, free text, grounds pleaded and other rhetorical role categories in the judgment. Within the single-judgment view the user-experience could be improved by incorporating hyperlinks that jump to a specific part of the judgment such as the target of the action or a specific citation. Moreover, utilising linked data to allow for contextual queries and complex visualisations of legal information is a useful tool to provide a better understanding of the role of a single judgment with respect to the entire judgment landscape.

6. Conclusions

There is a large amount of research in to the use of machine learning techniques on legal documents, whether for the goal of summarising judgments, identifying rhetorical roles, or

extracting entities. Often it is suggested that these methods could help increase accessibility of judgments, levelling the field between legal experts and the average citizen. However, there has been little work done to demonstrate how this might be possible.

In this paper we explained how the UK's new judgment publishing pipeline Find Case Law was developed, including the core annotators and the use of LegalDocML. We demonstrated how the work done to develop *Find Case Law* can be exposed to the end-user with little development effort with our system *MyJudgments*, which has been made publicly available on Github [7]. We also reviewed existing machine learning research undertaken on UK judgment data, demonstrating where the current state of UK legal ML research. We suggest that collaboration between researchers and public case law publication platform providers would improve the access to justice gap in the UK. Integrating various lines of machine learning research into a case law publication platform would provide greater insight into judgments in the UK, ensuring that citizens have clearer transparency of the inner workings of the judicial system.

References

[1] The European Commission, The 2018 EU Justice Scoreboard, The European Commission, 2018. URL: https://ec.europa.eu/info/sites/default/files/justice_scoreboard_2018_en.pdf.

[2] D. Hoadley, J. Tomlinson, E. Nemsic, C. Somers-Joce, How public is public law? the current state of open access to administrative court judgments, Judicial Review (2022). URL: https://doi.org/10.1080/10854681.2022.2111966.

[3] D. Hoadley, Part 2: Open access to english case law (the gaps), http://carrefax.com/articles-blog/2018/5/9/part-2-open-access-to-english-case-law-knackered-plumbing, 2018. Accessed: 2022-08-10.

[4] D. N. Byrom, Digital Justice: HMCTS data strategy and delivering access to justice, The Legal Education Foundation, 2019. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/835778/DigitalJusticeFINAL.PDF.

[5] Boost for open justice as court judgments get new home, https://www.gov.uk/government/news/boost-for-open-justice-as-court-judgments-get-new-home, 2021. Accessed: 2022-07-17.

[6] E. Filtz, S. Kirrane, A. Polleres, The linked legal data landscape: linking legal data across different countries, Artificial Intelligence and Law 29 (2021) 485–539.

[7] M. P. Fabio Vitali, V. Parisse, Akoma ntoso naming convention version 1.0, 2019. URL: http://docs.oasis-open.org/legaldocml/akn-nc/v1.0/akn-nc-v1.0.html, last accessed 8 August 2022.

[8] B. Hachey, C. Grover, Extractive summarisation of legal texts, Artificial Intelligence and Law 14 (2006) 305–345.

[9] O. Ray, A. Conroy, R. Imansyah, Summarisation with majority opinion., in: JURIX, 2020, pp. 247–250.

[10] D. Locke, G. Zuccon, Towards automatically classifying case law citation treatment using neural networks, in: Proceedings of the 24th Australasian Document Computing Symposium, 2019, pp. 1–8.

[7] https://github.com/mdrresearch/MyJudgments

[11] A. Kanapala, S. Pal, R. Pamula, Text summarization from legal documents: a survey, Artificial Intelligence Review 51 (2019) 371–402.
[12] J. Valvoda, O. Ray, K. Satoh, Using agreement statements to identify majority opinion in ukhl case law, in: Legal Knowledge and Information Systems, IOS Press, 2018, pp. 141–150.
[13] J. S. T. Howe, L. H. Khang, I. E. Chai, Legal area classification: A comparative study of text classifiers on singapore supreme court judgments, arXiv preprint arXiv:1904.06470 (2019).
[14] J. Gibson, R. Patel, C. Paskell, C. Peto, Justice data matters: Building a public mandate for court data use, https://research.thelegaleducationfoundation.org/wp-content/uploads/2022/07/Justice-Data-Matters-Report-Final-.pdf, 2022. Accessed: 2022-08-09.
[15] M. Honnibal, I. Montani, S. Van Landeghem, A. Boyd, spacy: Industrial-strength natural language processing in python (2020). URL: https://doi.org/10.5281/zenodo.1212303.
[16] M. Neumann, D. King, I. Beltagy, W. Ammar, ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing, in: Proceedings of the 18th BioNLP Workshop and Shared Task, Association for Computational Linguistics, Florence, Italy, 2019, pp. 319–327. URL: https://www.aclweb.org/anthology/W19-5034. doi:10.18653/v1/W19-5034. arXiv:arXiv:1902.07669.
[17] B. Hachey, C. Grover, A rhetorical status classifier for legal text summarisation, in: Text Summarization Branches Out, 2004, pp. 35–42.
[18] P. Bhattacharya, S. Paul, K. Ghosh, S. Ghosh, A. Wyner, Deeprhole: deep learning for rhetorical role labeling of sentences in legal case documents, Artificial Intelligence and Law (2021) 1–38.