

April 2020

An explainable approach to deducing outcomes in European Court of Human Rights cases using ADFs

Joe COLLENETTE, Katie ATKINSON and Trevor BENCH-CAPON
Department of Computer Science, University of Liverpool, UK

Abstract. In this paper we present an argumentation-based approach to representing and reasoning about a domain of law that has previously been addressed through a machine learning approach. The domain concerns cases that all fall within the remit of a specific Article within the European Court of Human Rights. We perform a comparison between the approaches, based on two criteria: ability of the model to accurately replicate the decision that was made in the real life legal cases within the particular domain, and the quality of the explanation provided by the models. Our initial results show that the system based on the argumentation approach improves on the machine learning results in terms of accuracy, and can explain its outcomes in terms of the issue on which the case turned, and the factors that were crucial in arriving at the conclusion.

Keywords. Legal case-based reasoning, Abstract Dialectical Frameworks, Explanation

1. Introduction

Reasoning with legal cases has always been a central topic of AI and Law. The basic feature of this sort of reasoning is that there is a body of case law - previous decisions - and a new case must be decided in the light of these precedents. A hearing of a case takes the form of each side presenting arguments, typically based on these precedent decisions, as to why they should win¹.

The adversarial nature of law naturally led to such reasoning being seen in terms of argumentation, with arguments being presented for both sides and the user being expected to choose between them. Examples of such systems are McCarty's TAXMAN [3], HYPO [4] and the many systems influenced by HYPO [5]. Advantages of this approach were that it was able to offer a model of legal reasoning, and provide a full explanation of the reasoning. Early systems only presented the arguments without offering a decision, but it is also possible to use such systems for prediction by attempting to assess the com-

¹Although precedent cases are more especially associated with the common law traditions of UK and USA, precedents also play a role in European civil law, guiding the permissible interpretations. The role of precedents in Civil Law is discussed in [1], who quotes [2]: "precedent now plays a significant part in legal decision making and the development of law in all the countries and legal traditions that we have reviewed," including nine civil law jurisdictions (e.g., Germany, France, Italy, and Spain) and two common law jurisdictions (the United Kingdom and New York State)".

April 2020

peting arguments. An example from the rule based tradition is [6] and an example from the case based tradition is [7]. The argumentation approach continues to be popular: recent work includes [8], [9] and [10] from ICAIL 2019 and a commercial implementation [11].

Although this approach has proven successful in modelling aspects of legal reasoning and providing outcomes backed with full explanations that are readily understandable in legal terms, it can require a good deal of expertise to analyse the precedents and construct a model of the case law, and so encounters the classic *knowledge engineering bottleneck*. Given that the problem is essentially classifying new examples on the basis of a large amount of previous cases, it would seem that the use of machine learning techniques might enable this bottleneck to be avoided. Therefore even several decades ago efforts were made to apply machine learning techniques to case law, such as [12], and [13]. This did produce some apparently promising results, while attracting some theoretical criticism [14]. But the practical difficulties were perhaps even more decisive: at that time the machine learning techniques could not be applied to natural language, and so analysis was needed to identify the relevant features of the cases and to ascribe them to the training examples, and the amount of case data available in suitable form in those pre-internet days was limited. For these reasons these approaches did not become mainstream.

Recently, however, because of the improvement in machine learning techniques and the vastly increased availability of data, there has been a marked revival of interest in these techniques. There is an annual Competition on Legal Information Extraction/Entailment (COLIEE)², the sixth edition of which was held in 2019. Some projects have attracted considerable attention, perhaps most notably [15], which attracted a great deal of media interest³.

We may take [15] as representative of this new trend in using learning from large data sets to predict case outcomes. It was designed to classify cases heard before the European Court of Human Rights into violations and non-violations, according to three Articles of the European Convention on Human Rights. It was claimed that “our models can predict the court’s decisions with a strong accuracy (79% on average)”. The study used a dataset comprising 250 cases related to Article 3, 80 to Article 6 and 254 to Article 8, balanced between violation and non-violation cases. Textual information was represented using contiguous word sequences, i.e., N-grams, and topics, and used to train Support Vector Machine (SVM) classifiers [16]. The same domain was also examined in [17] where broadly similar results were obtained using more cases, with a slightly lower success rate (75% on average). That paper, however, reported another experiment, designed to test the usefulness for predicting future cases on the basis of data about the past. In this experiment they trained on data up to 2013 and tested on data from 2014-5 and from 2016-7. This showed a decrease in performance. For example, consider Article 6: whereas using all the data without regard to time produced a success rate of 80%, this fell to 64% for the 2014-5 data and to 63% and to only 59% for 2016-7 cases when the training set is all before 2005.

²<https://sites.ualberta.ca/~rabelo/COLIEE2019/>

³See for example, Artificial intelligence ‘judge’ developed by UCL computer scientists, *Guardian*, October, 2016. <https://www.theguardian.com/technology/2016/oct/24/artificial-intelligence-judge-university-college-london-computer-scientists>.

April 2020

This experiment bears out an important criticism made of the machine learning for prediction approach in works such as [18], that because case law evolves over time, changes may not be picked up by a system trained on historic data. Another key criticism is the success rate: although [15] describes 79% as “strong accuracy”, having over a fifth of cases decided wrongly would not be acceptable in a legal system. Finally there is the lack of explanation: in law the parties to the suit have a *right* to explanation [19]. Moreover this explanation should be couched in legally relevant terms. No explanation is provided in [15], beyond a list of 20 most frequent words, listed in order of their SVM weight. These lists do not really provide a satisfying explanation, typically containing names of months and common words such as “court, applicant, article, judgment, case, law, proceeding, application, government”, which head the list for topic 23 of Article 6.

For these reasons we decided to tackle the domain of [15] using argumentation based methods. Our aim is to show that, using proportionate effort, a system can be built with these methods which is more accurate, provides adequate explanation facilities and is structured so as to be amenable to changes in the case law [20]. Unlike machine learning approaches, however, this system will require users to have sufficient legal expertise to answer questions about the case. Our chosen approach is to use the ANGELIC methodology [21], of which we will give a brief overview in the next section.

2. Background: the ANGELIC methodology

The ANGELIC methodology was designed to encapsulate the knowledge of a body of case law in a way which would support argumentation and facilitate future modification as that case law evolved. It is described in [21] and was used to support the development of applications in [11] and [22]. Over the years, representing cases in terms of factors had become the *de facto* standard approach to reasoning with legal cases. These factors can be usefully organised into a factor hierarchy in the manner of [23] to show how they relate to issues in the domain. In a factor hierarchy the children of a node are reasons for the presence or absence of their parent. In ANGELIC this factor hierarchy is interpreted as an Abstract Dialectical Framework (ADF) [24]. In an ADF each node has acceptance conditions which specify the conditions under which a node will be accepted or rejected in terms of its children. The factor hierarchy and the ADF have the same structure, nodes with positive and negative children, but the ADF enables the relation between a parent and its children to be specified precisely. In ANGELIC the acceptance conditions take the form of a prioritised set of conditions for the acceptance or rejection of the node, each individually sufficient, and collectively necessary. The effect is that once the values of the children have been established, we can use the acceptance conditions to provide arguments to accept or reject the parent. In turn, as we descend the tree, we get arguments to accept and reject the children. The leaves of the tree take the form of questions, to be answered from the facts of the case. This produces the kind of argument-subargument structure found in ASPIC+ [25]. The relation between the ADFs of ANGELIC and ASPIC+ is discussed in [26]. Because the acceptability of the parent depends only on the acceptability of its children, we get a highly modular structure to support maintenance and to accommodate changes [20].

The ADF produced using ANGELIC is always a tree and the nature of the nodes changes as we descend the tree, as described in [27]. The root of the tree is the *verdict*,

April 2020

the overall question to be decided, such as whether there is a violation. The next layers represent *issues*, the various broad ways in which an answer to the question must be considered, such as the various ways in which an Article may be violated. Often, the issues can be found in the legislation. Below these are the *abstract factors*, the various considerations found relevant to the issues in previous cases. Below these are the *base level factors*, which are the legal facts as accepted by the court. These in turn unfold into the plain facts of the case, so that the base level factors can be resolved by posing questions to the user.

Thus the ADF produced by ANGELIC provides a systematic and modular representation of the case law, from which an argument from case facts to issues and verdict can readily be recovered.

3. Domain application: Cases on Article 6 in the European Court of Human Rights

In this section we describe the ADF produced for Article 6, in addition to providing a brief summary of Article 6. The previous work by Aletras et al. [15] used 80 cases that relate to Article 6: throughout this paper we will use some of these cases as examples to highlight how we have developed and implemented our ADF representation of the domain.

Article 6 of the European Convention on Human Rights guarantees the right to a fair trial. The aim of the article is to guarantee the procedural rights of parties in civil proceedings and the rights of defendants in criminal proceedings. In essence the article is concerned with whether an applicant had ample opportunity to present their side of the case and contest any aspects they consider to be false, rather than ensuring that the courts have come to the correct decision. Provided the procedures followed were acceptable, the decision is acceptable with respect to Article 6.

When developing the ADF which represents Article 6, the *verdict* is whether there has been a violation of Article 6. The *issues* which inform whether there has been violation, these were determined from the legislation. There are three substantive issues that come from the legislation, and there are two additional procedural issues that need to be considered before examining the substantive issues. The three substantive issues, which have been heavily summarised, are:

- The case was fair and public
- The applicant was presumed innocent until proven guilty
- The applicant had their minimum rights respected

The two additional procedural issues are:

- The applicant bringing the case is the victim who was the discussed case
- The case is admissible

These five main issues are decided by considering *abstract factors*. These abstract factors may in turn have further abstract factors which decide their validity. In total our developed ADF has thirty-five abstract factors. For example when considering the third main issue, that the applicant had the minimum rights, there are six abstract factors which have been determined from the legislation and which decide that issue for the given case. These presence of these abstract factors themselves is decided by further abstract factors.

April 2020

We then come to the bottom of the ADF where we reach the *base level factors*: these are answers to questions given by the user on the basis of the facts of the case.

A full example, where we follow a particular branch to its base level factor, is as follows:

Verdict There has been a violation of Article 6

This verdict relates, as stated above, to five issues, one of which is:

Issue The case is admissible

This issue is determined by two abstract factors:

Abstract Factor The case is well-founded

Abstract Factor There was no significant disadvantage

If either of these are not acceptable, the case will fail on the issue. That *there was no significant disadvantage* descends into further abstract factors, but that *the case is well founded* can be accepted on the basis of three base level factors, relating to the following three questions:

Base Level Factor Question Has the case been trivially answered previously?

Base Level Factor Question Does the applicant have evidence which supports the breach of Article 6?

Base Level Factor Question Is the case nonsensical?

If the answer is *yes* to the first two, and *no* to the third, then the case can be accepted as admissible and the substantive issues considered.

From an examination of the application document lodged by the claimant it is possible for a person familiar with such documents to determine that there is some evidence to back the claim, and that the claim is not nonsensical. The answers to these questions do require some familiarity with existing case law, but this can be expected from a lawyer using the system.

Table 1 contains a small subset of all the issues and factors that we have developed for the ADF. The complete ADF contains 51 nodes: 1 verdict, 5 issues, 10 abstract factors and 35 base level factors, which can be ascribed on the basis of an answer to a corresponding question. The questions relevant to the nodes in Table 1 are given in Table 2.

3.1. Implementation in Prolog

To implement the above ADF as an executable computational program, we followed a similar path to Al-Abdulkarim et al. [27] where each case is represented by a list of base level factors in a Prolog program. The program will print the status of each factor as it is determined, by resolving the ADF structure in the program.

The Prolog program follows the European Court by firstly checking if the case is admissible. Only once the case had been declared admissible, the program will traverse the full ADF and report the findings it produces. The code snippet below shows how node ID 1 from Table 1 has been developed in the Prolog code. The program resolves what the values of *X*, *Y* and *Z* are before checking the conditions as described in the table before printing human-readable output. Note that the identity tests are required to ensure that every node is visited and so can be included in the explanation.

April 2020

Table 1. Subset of issues and factors in the ADF describing Article 6

ID & Factor	Children	Conditions
1 - Violation of Article 6	[2,3,8,20,21]	REJECT IF NOT is a victim OR NOT case is admissible ACCEPT IF the case was not fair or public OR victim was presumed guilty OR the victim did not have the minimum rights REJECT otherwise
2 - Is a victim		REJECT IF Q2 ACCEPT otherwise
3 - The case is admissible	[4,5]	REJECT IF NOT the case is well-founded OR there was no significant disadvantage ACCEPT otherwise
4 - The case is well founded		ACCEPT IF Q4a OR Q4b OR NOT Q4c REJECT otherwise
5 - No significant disadvantage	[6,7]	REJECT IF there is no fundamental reason why the case should be looked at OR there are domestic tribunals that have not looked at the case ACCEPT otherwise

Table 2. Subset of base level factors which answer the leaf node abstract base factors (LN)

Q	Base level factor question
2	Was the person bringing the case the victim?
4a	Has the case been trivially answered previously?
4b	Does the applicant have evidence which supports the breach of Article 6?
4c	Is the case nonsensical?

```
violationOfArticle6(CASE) :- (isVictim(CASE, X),
    isAdmissible(CASE, Y)), X == valid, Y == valid,
    fullcheck(CASE), !.
violationOfArticle6(_) :-
    write("The case is therefore inadmissible"), nl.

fullcheck(CASE) :- (isFairAndPublic(CASE, X),
    isPresumedInnocent(CASE, Y), hadMinimumRights(CASE, Z)),
    (X == valid, Y == valid, Z == valid),
    write("Therefore there is no violation of Article 6")
    , nl, !.
fullcheck(_) :-
    write("Therefore there is a violation of Article 6"), nl.
```

To resolve what the value of X is in the example, the program will continue checking conditions and printing output in order to give the value to X that has been requested.

April 2020

For example *isFairAndPublic* returns valid when the conditions of the ADF have been met and returns invalid otherwise, which is shown clearly in the code snippet below.

```
isFairAndPublic(case(_,L), valid) :-
    (isConductedInAReasonableTime(case(_,L),X),
     isIndependantAndImpartial(case(_,L), Y),
     isConductedPublicly(case(_,L), Z),
     isEqualityOfArms(case(_,L), A),
     givenAccessToCourt(case(_,L), B)),
    (A == valid, B == valid, X == valid,
     Y == valid, Z == valid),
    write("The case was fair and public"),
    nl, !.
isFairAndPublic(case(_,_), invalid) :-
    write("The case was was not fair and public"), nl.
```

The program will continue traversing the ADF until it reaches factors that can be resolved by checking the answers to a base level factor question. For example Q2, answers whether the applicant is the victim in question. If the question should be answered positively in the case, then it is included in the list that is provided at the start of the program (*L* is the code fragments), as shown in the code snippet below.

```
isVictim(case(_,L), valid) :- member(Q2, L),
    write("The applicant is the victim"), nl.
isVictim(case(_,L), invalid) :- not(member(Q2,L)),
    write("The applicant is not the victim"), nl.
```

These three aspects (*fullCheck* for the issues, a procedure for each abstract factor, and the test against the list of base level factors) make up the entire Prolog program and when run will produce the list of factors that determine the decision of the program.

4. Experiments

For our implemented Prolog program, we have manually ascribed base level factors to 10 different cases to test different parts of the Prolog program. The cases chosen are all from the European Court of Human Rights and require resolution regarding whether there was a violation of Article 6.

April 2020

Table 3. Cases used to test the Prolog program, along with highlights of the output produced by the Prolog program

Case	Actual outcome	Highlights of program output
MARGUŠ v. CROATIA	No violation	Therefore there is no violation of Article 6.
CARDOT v. FRANCE	Inadmissible	Not all domestic courts have been exhausted The applicant suffered a disadvantage The case is therefore inadmissible
ABDULLAYEV v. RUSSIA	Violation	Not given appropriate access to a court The case was not fair and public The applicant has not waived right to defend themselves Not prevented from accessing lawyers The applicant is defending themselves in person Therefore there is a violation of Article 6
ZARKOV v. SERBIA	Violation	The Government caused unreasonable delays The case was not conducted in a reasonable time Therefore there is a violation of Article 6
MOSER v. AUSTRIA	Violation	The case was not pronounced publicly The case was not conducted publicly The case was required to be conducted publicly, and was not There was not an equality of arms Therefore there is a violation of Article 6
CHAPMAN v. THE UNITED KINGDOM	No violation	Therefore there is no violation of Article 6
KHANUSTARANOV v. RUSSIA	Violation	Not given appropriate access to a court Therefore there is a violation of Article 6
STOILKOVSKA v. THE FORMER YUGOSLAV REPUBLIC OF MACEDONIA	Violation	The Government caused unreasonable delays The case was not conducted in a reasonable time Not given appropriate access to a court The case was not fair and public Therefore there is a violation of Article 6
UŽKURĖLIENĖ AND OTHERS v. LITHUANIA	No violation	Therefore there is no violation of Article 6
T.P. AND K.M. v. THE UNITED KINGDOM	No violation	Therefore there is no violation of Article 6

April 2020

To illustrate how the cases have been used to test our Prolog program, we will walk through Moser v. Austria⁴. We have analysed each case and manually identified the factors that will be input into the program, answering each of the base level questions. For example, question 2 (Q2) holds as Moser was the victim in the case that is being debated. Q19 and Q20 do not hold as the verdict was not pronounced or conducted publicly. Q21 also does not hold as the victim was unable to comment on reports that were used by the Austrian courts. While a number of the questions are easily answered, such as Q2, answers to other questions, such as Q25, need to be derived from assumptions. For example for Q25 (“Was the victim informed of the crime in a language they understand?”), there is no discussion of this in the case facts, therefore we assume that the victim was told in a language they understand and Q25 should be considered true. The program, as with the European Court, assumes there was no violation unless there is specific discussion stating the reasoning why there is a violation. We believe that had the victim not been informed in an appropriate language, that would have appeared in the application.

Table 3 shows the results for all ten cases that we have chosen, each case has been chosen to test a specific aspect of the program. We have chosen nine cases that were used as part of the Article 6 dataset in Aletras et al. [15]. In addition we have added Cardot v. France⁵ in order to test a case that is inadmissible. Note that Aletras et al. [15] used post trial documents and so did not include any inadmissible cases. We believe that it is important for an implementation to assess whether the case is admissible and so demonstrate this through our test set.

From the results we can see that the implementation of the ADF achieves correct results for all 10 test cases, and with the ability to explain why the violation was reported, or was not reported. Whilst our sample size for this initial experiment is small, the cases have been carefully selected to exercise different branches of the program and hence are extremely encouraging. We believe that when compared to the Aletras et al. [15] approach, there are a number of benefits to our approach which will help with finding acceptance among lawyers interested to use these programs to provide decision-support in case management, as has been shown through [11]. Next follows the output the program produces from Moser v. Austria. Issues are given in italics and the factors that led to the violation are indicated by “***”. It is these factors which provide the highlights column in Table 3.

The applicant is the victim

The case is well founded

The case does examine a fundamental part of human rights act

Domestic courts have been exhausted

The applicant suffered no disadvantage

The case is therefore admissible

The Government did not cause unreasonable delays

The case was conducted in a reasonable time

The Government was subjectively impartial

The Government was objectively impartial

The Government was independent and impartial

Public hearing would not prejudice outcome of case

⁴MOSER v. AUSTRIA JUDGMENT, 2006, European Court of Human Rights, (Application no. 12643/02)

⁵CARDOT v. FRANCE JUDGMENT, 1991, European Court of Human Rights, (Application no. 11069/84)

April 2020

Safety of the public would not be impacted by the case being-
publicly pronounced
Privacy is not required to deliver justice
The public hearing would not hinder delivery of justice
***The case was not pronounced publicly
***The case was not conducted publicly
***The case was required to be conducted publicly, and was not
**There was not an equality of arms
Given appropriate access to a court
****The case was not fair and public*
The Government bore the burden of proof
Any doubt benefited the applicant
The applicant was presumed innocent until guilty
Was informed in the correct language
Was promptly detailed circumstances to mount a reasonable-
defence
Applicant was told what crime they had committed
The applicant was informed promptly in a language they understand
Did have time or facilities to mound a reasonable defence
The applicant has not attempted to escape trial
The applicant has not waived right to defend themselves
Not prevented from accessing lawyers
The applicant is defending themselves in person
Free access to legal assistance was available
The applicant therefore had access to legal assistance
Any witnesses were examined under the same different-
conditions when compared to the Government
Any witnesses that were not present had valid reasoning
The applicant therefore was able to examine witnesses
Had access to interpreter as required
The applicant had the minimum rights required
***Therefore there is a violation of Article 6

The biggest strength of the ADF is that it is able to explain through a series of statements the line(s) of reasoning which produces the outcome for a case, and so give a complete explanation in terms used in the domain. The output could be post-processed as in [21] to improve the quality of the presentation.

Although the series of test cases does suggest that the ADF can deduce case outcomes with an accuracy better than than the 79% produced by [15], future experiments on a larger test set are needed to further confirm this. The reason why our approach leads to better results is that it is based on an understanding of the domain, rather than the machine learning approach which creates learned probabilities on word groupings lacking full context, which is not how court cases are decided in practice.

Even if a program produced by machine learning was a perfect predictor of Article 6 cases, the program will over time become worse at predicting the outcomes due to how the law and its interpretation changes over time [17]. Programs that predict the outcomes of court cases must be able to adapt quickly to new information in order to capture the

April 2020

change in interpretation. Such changes are typically produced by landmark cases, which signal that updating is required. An ADF approach, such as the one developed, can be adapted quickly by changing the questions and factors, exploiting its modular nature. Adapting a machine learning program when a landmark case arises would be far more problematic. All of the past cases would be called into question: although many would be unaffected by the new ruling, some will now give a misleading picture of the law. Either all cases must be analysed to identify those rendered ineffective, which would remove many of the advantages of such systems, or there has to be time to acquire a substantial training set reflecting the new understanding of the law.

While there are several benefits to our approach, the major drawback is the time and expertise required to ascribe the factors that are fed into the program to describe a new case. We believe that this is the point at which learning from texts can produce benefits by answering the questions instantiating the base level factors, similar to the approaches by Ashley and Brüninghaus [28] and Branting et al. [29]. The aim of employing machine learning in this way is that we keep the benefits of good old-fashioned AI (through the ADF), which is needed in order to satisfy the lawyers who demand extremely high accuracy and explainability, while speeding up and reducing the expertise needed for processing new individual cases.

5. Conclusion

We have presented an argumentation-based representation of Article 6 of the European Convention on Human Rights, which is the right to a fair trial. The representation is an Abstract Dialectical Framework produced using the ANGELIC methodology, where the domain is represented as a tree, with the root node being a verdict, followed by children which are issues, whose children are abstract factors. The leaf nodes are base level factors, which are ascribed to a case by answering questions about the facts of the case. This framework was implemented in Prolog to enable reasoning about cases within the domain. Our framework and the Prolog program were tested with cases that concern Article 6 (and perhaps other articles) and were resolved in the ECHR.

Whilst the exercise in itself has been an instructive demonstration of the application of a computational model of argument to a real world domain, of further interest is that we were able to compare our approach to a machine learning approach to predicting cases in the exact same domain. Our success in producing a high level of accuracy in the performance of the program, being able to readily adapt the program as the law evolves, *and* being able to accompany this with strong explanatory features, addresses several limitations associated with machine learning approaches.

References

- [1] Ashley KD. Case-based models of legal reasoning in a civil law context. In: International congress of comparative cultures and legal systems of the instituto de investigaciones jurídicas. Universidad Nacional Autónoma de México, Mexico City; 2004. p. 1–30.
- [2] MacCormick DN, Summers RS, Goodhart AL. Interpreting precedents: a comparative study. Routledge; 2016.
- [3] McCarty LT. Reflections on TAXMAN: An experiment in Artificial Intelligence and legal reasoning. Harvard Law Review. 1976;90:837.

April 2020

- [4] Rissland EL, Ashley KD. A case-based system for Trade Secrets law. In: Proceedings of the 1st International Conference on Artificial Intelligence and Law. ACM; 1987. p. 60–66.
- [5] Bench-Capon T. HYPO's legacy: introduction to the virtual special issue. *Artificial Intelligence and Law*. 2017;25(2):1–46.
- [6] Prakken H. A tool in modelling disagreement in law: preferring the most specific argument. In: Proceedings of the 3rd international conference on Artificial intelligence and law; 1991. p. 165–174.
- [7] Brüninghaus S, Ashley KD. Predicting outcomes of case based legal arguments. In: Proceedings of the 9th International conference on Artificial Intelligence and Law. ACM; 2003. p. 233–242.
- [8] Atkinson K, Bench-Capon T. Reasoning with Legal Cases: Analogy or Rule Application? In: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law; 2019. p. 12–21.
- [9] Prakken H. Modelling accrual of arguments in ASPIC+. In: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law; 2019. p. 103–112.
- [10] Westermann H, Walker VR, Ashley KD, Benyekhlef K. Using Factors to Predict and Analyze Landlord-Tenant Decisions to Increase Access to Justice. In: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law; 2019. p. 133–142.
- [11] Al-Abdulkarim L, Atkinson K, Bench-Capon T, Whittle S, Williams R, Wolfenden C. Noise induced hearing loss: Building an application using the ANGELIC methodology. *Argument & Computation*. 2019;10(1):5–22.
- [12] Philipps L. A Neural Network to Identify Legal Precedents. In: Proc. of the 9th Symposium on Legal Data Processing in Europe. Council of Europe, Publishing and Documentation Service; 1089. p. 99–106.
- [13] Bench-Capon T. Neural networks and open texture. In: Proceedings of the 4th international conference on Artificial intelligence and law; 1993. p. 292–297.
- [14] Hunter D. Looking for law in all the wrong places: Legal theory and legal neural networks. In: Proceedings of JURIX 1994; 1994. p. 55–64.
- [15] Aletras N, Tsarapatsanis D, Preoțiuc-Pietro D, Lampos V. Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science*. 2016;2:e93.
- [16] Vapnik V. The nature of statistical learning theory. Springer science & business media; 2013.
- [17] Medvedeva M, Vols M, Wieling M. Using machine learning to predict decisions of the European Court of Human Rights. *Artificial Intelligence and Law*. Available on-line, 2019;p. 1–30.
- [18] Bench-Capon T. The Need for Good Old fashioned AI and Law. In: *International Trends in Legal Informatics: A Festschrift for Erich Schweighofer*. Editions Weblaw, Bern; 2020. p. 23–36.
- [19] Doshi-Velez F, Kortz M, Budish R, Bavitz C, Gershman S, O'Brien D, et al. Accountability of AI under the law: The role of explanation. arXiv preprint arXiv:171101134. 2017;.
- [20] Al-Abdulkarim L, Atkinson K, Bench-Capon T. Accommodating change. *Artificial Intelligence and Law*. 2016;24(4):409–427.
- [21] Al-Abdulkarim L, Atkinson K, Bench-Capon T. A methodology for designing systems to reason with legal cases using ADFs. *Artificial Intelligence and Law*. 2016;24(1):1–49.
- [22] Atkinson K, Bench-Capon T, Routen T, Sánchez A, Whittle S, Williams R, et al. Realising ANGELIC Designs Using Logiak. In: *Legal Knowledge and Information Systems: JURIX 2019: The Thirty-second Annual Conference*. vol. 322. IOS Press; 2019. p. 151.
- [23] Alevin V. Teaching case-based argumentation through a model and examples. University of Pittsburgh; 1997.
- [24] Brewka G, Woltran S. Abstract dialectical frameworks. In: *Twelfth International Conference on the Principles of Knowledge Representation and Reasoning*; 2010. p. 102–111.
- [25] Prakken H. An abstract framework for argumentation with structured arguments. *Argument and Computation*. 2010;1(2):93–124.
- [26] Atkinson K, Bench-Capon T. Relating the ANGELIC Methodology and ASPIC+. In: *COMMA*; 2018. p. 109–116.
- [27] Al-Abdulkarim L, Bench-Capon T, Atkinson K. Statement types in legal argument. In: *Legal Knowledge and Information Systems: JURIX 2016: The 29th Annual Conference*. vol. 294. IOS Press; 2016. p. 3.
- [28] Ashley KD, Brüninghaus S. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*. 2009;17(2):125–165.
- [29] Branting K, Weiss B, Brown B, Pfeifer C, Chakraborty A, Ferro L, et al. Semi-Supervised Methods for Explainable Legal Prediction. In: *Proc. of the 17th International Conf. on AI and Law*; 2019. p. 22–31.