Explainable AI Tools for Legal Reasoning about Cases: A Study on The European Court of Human Rights

Joe Collenette

Department of Computer Science, University of Manchester, UK

Katie Atkinson, Trevor Bench-Capon

Department of Computer Science, University of Liverpool, UK

Abstract

In this paper we report on a significant research project undertaken to design, implement and evaluate explainable decision-support tools for deciding legal cases. We provide a model of a legal domain, Article 6 of the European Convention on Human Rights, constructed using a methodology from the field of computational models of argument. We describe how the formal model has been developed, extended and transformed into practical tools, which were then used in evaluation exercises to determine the effectiveness and usability of the tools. The underpinning AI techniques used yield a level of explanation that is firmly grounded in legal reasoning and is also digestible by the target end users, as demonstrated through our evaluation activities. The results of our experimental evaluation show that on the first pass, our tool achieved an accuracy rate of 97% in matching the actual decisions of the cases and the user studies conducted gave highly encouraging results with respect to usability. As such, our project demonstrates how trustworthy AI tools can be built for a real world legal domain where critical needs of the end users are accounted for.

Keywords: Computational models of argument, Legal reasoning, Explainable AI, European Convention on Human Rights

1. Introduction

Modelling and supporting legal decision making and predicting the outcome of legal cases have been central topics of AI and Law since its beginnings in the 1970s [1]. Typically the aim is not to replace lawyers, but to provide support by identifying issues and indicating the likely consequences of the facts of a case. Many approaches have been developed over the last five decades, perhaps most notably those deriving from the HYPO project of Rissland and Ashley [2], [3]. In recent years, the topic has attracted much interest arising from the increasing use of machine learning to perform the task. For example, a number of projects have addressed the prediction task in the context of the European Convention of Human Rights (ECHR) including [4], [5], [6], [7] and [8]. These studies all report success[1]. There are, however, reasons to believe that machine learning approaches have only a limited role in supporting legal decisions, centring around the lack of explanation, the difficulties in adapting to changes in the law, and the possibility of historic bias being implicit in the data, which would go unnoticed without explanations being provided [9]. Further limitations on the use of these algorithms to support legal decision making are identified in [10] and [11]. Moreover there are questions relating to the information that is available before the trial [12]: the systems referred to above

[1] Predictions are generally reported as correct in around 70-85% of cases, using historic data for both training and test sets. JURI Says, the program described in [8], originally reported a success rate in this range, but has since been applied to new cases as they are decided and over time this accuracy has fallen. The accuracy for February 2022 was 76.9%, but it was only 55.9% over the last year. Its monthly figure fluctuates greatly: it was 87.2% for January 2021 but fell to 48.9% for November 2021. The overall accuracy since it has been running is currently 59.1%. JURI Says can be found at `https://jurisays.com/` (accessed 2022/07/26). It is, however, doubtful whether accuracy of less than 90% would be acceptable for practical deployment in this application. The deterioration of the performance over time is a further indication of the problems inherent in using historic data to train a system to predict future cases in a domain such as law which is subject to constant change, not only when legislation is revised but also as case law evolves to reflect changing social attitudes. These problems were already suggested by the second experiment in [5].

obtain the case facts from the decisions which are written after the decision has been made, in order to justify that decision, and so may to some extent anticipate the outcome. These limitations suggest that there will be a continuing role for more traditional knowledge representation techniques when building legal systems, perhaps in partnership with machine learning systems [13], [14] and [15].

An important feature of legal applications is the centrality and indispensability of explanations. In a legal proceedings, participants have a right to an explanation of the decision in their case [16], to persuade the loser that the decision was correct, or to form the basis of an appeal, and to enable public scrutiny of the verdict. Without an explanation, a bare decision offers no support to a judge [10], and provides potential litigants with no assistance in presenting their case. In consequence, explanation has always been at the centre of traditional AI and Law systems investigating these issues [17]. In contrast, those prediction systems based on machine learning approaches do not provide satisfactory explanations. It has been proposed that explanation techniques developed in more traditional systems could be used to explain the output of prediction algorithms (e.g. [14] and [15]), but this will require an underlying domain model. In this paper we will describe research undertaken to design, implement and evaluate an explainable decision-support tool for deciding legal cases under Article 6 of the European Convention on Human Rights (ECHR). The ECHR domain was chosen for our case study because of the ready availability of case materials in the public domain, and because this domain has been the target for a number of different machine learning based approaches, enabling comparison with them.

In earlier work [17] we provided a comprehensive survey of the landscape of techniques for explanation in AI and Law, and identified paths for future developments based on gaps yet to be addressed. In this paper, we have progressed one such development strand to demonstrate how expert knowledge within a domain of law can be captured in order to automate reasoning about legal cases and provide explanations of outcomes that are easily digestible by the legal professionals at whom the tools are aimed. The key new contributions of the work

reported in this paper can be summarised as:

- Production of a novel, legally-grounded symbolic model of a complex real world legal domain, achieved through application of a methodology using techniques from computational models of argument;
- New practical tools aimed at end users to enable them to undertake processing of legal work through AI-based support that provides a high level of explainability, going beyond the current state-of-the-art;
- Results from three evaluation exercises, including studies involving real world end users, accompanied by an analysis on the viability and usability of our research and its application in a real world legal setting;
- A demonstration of how explainable AI can be developed for a real world problem such that explanations provided by the tools are presented using terms and a structure that are familiar to domain users, thus promoting trustworthiness.

Section 2 will supply the background to the project by summarising the various approaches to modelling legal reasoning and supporting decisions in legal cases that have been developed in AI and Law. Particular emphasis will be placed on how the approaches attempt to meet the particular requirements of legal applications with respect to maintenance and explanation. Section 3 will describe the ECHR domain, in particular Article 6, which has been the focus of our work. Section 4 describes our representation of the domain as an Abstract Dialectical Framework [18], following the methodology proposed in [19]. Section 5 describes the implementation of this model and Section 6 the evaluation of this implementation. Finally, Section 7 offers some concluding remarks.

2. Background: Modelling Legal Reasoning

Before turning to our case study, in section 2.1 we give an overview of the knowledge that is needed to predict legal cases, and some essential requirements

on computer applications to support legal decision making that the knowledge representation will need to facilitate. We then review, in section 2.2, how knowledge representation for predicting cases has developed in the AI and Law literature. Note that the systems we describe are intended only to *support* legal decision making by providing reasoned explanations for case outcomes. It is, of course, accepted that not every nuance of legal reasoning will be captured, which is why it is important that a legally qualified user assess the arguments offered by the system, to ensure that no subtleties have been missed. Finally, in 2.3 we advance a proposal for using a contemporary knowledge representation technique, and describe how this has been used in practice.

2.1. Knowledge Required

There are two primary sources of law: legislation and cases. There are also a number of secondary sources such as commentaries, but these are concerned with how the law should be interpreted, and so inform the way the primary sources are represented, rather than being themselves represented.

2.1.1. Legislation

Legislation is typically presented as a set of definitions, or rules for the application of the legal concepts. Thus in Section 1 the UK Theft Act of 1968 we find

> (1) A person is guilty of theft if he dishonestly appropriates property belonging to another with the intention of permanently depriving the other of it;

This stipulates the conditions which must be satisfied if a person is to be found guilty of theft. But there are a number of terms which need to be interpreted, and these may become the subject of dispute. These terms may be defined further in the legislation, thus for the *Theft Act 1968*, "Dishonestly", "Appropriates", "Property", "Belonging to another", "With the intention of permanently depriving the other of it" are each defined in turn in sections 2-6

of the *Theft Act*. These definitions, however, themselves contain terms which stand in need of interpretation, and at some point the legislation will stop, and it will be the role of the courts to apply the law in the light of the particular circumstances of the cases brought before them. The knowledge of how to interpret these terms is found in the reported decisions[2] made in precedent cases.

2.1.2. Case Law

It is a fundamental principle of justice that like cases should be treated in a like manner. In order to achieve this, the person deciding a case must be aware of what was decided in similar cases in the past, and follow those decisions unless there is good reason not to do so. In Common Law traditions such as those of the UK and US, this principle is formalised in the doctrine of *stare decisis* ("let the decision stand") which obliges decisions of the appropriate status to be followed when deciding a new case. In Civil Law traditions such as are found in Europe, this element is lacking, but none the less, previous decisions are considered and typically respected [20], [21]. For the ECHR, which will be the subject of our case study, although the Court's previous judgments are not formally binding on the Court, it does not deviate from them without a very good reason and does so very rarely. The jurisprudence literature gives a number of models of precedential constraint. Those discussed in [22] include several which have been used in AI and Law, including: *balance of factors* (used in e.g. [23] and [24]), in which reasons for a party are weighed against reasons against that party, *purposive* (used in e.g. [25] and [26]), in which the decision is made so as to

[2]In the UK, about 2500 judgments (less than 2% of all judgments) are reported in law reports series each year. Decisions of the Supreme Court (previously House of Lords) and the Court of Appeal predominate because of the weight accorded them by the doctrine of precedent. Only a small proportion of the thousands of first instance cases in the High Court are reported (https://www.law.ox.ac.uk/legal-research-and-mooting-skills-programme/law-reports). This selective approach to decisions in previous cases contrasts with machine learning approaches which do not make an assessment of the importance of decisions.

promote the social purposes of the law concerned; and *rule based* (used in e.g. [27] and [28]), in which the precedent cases are seen as the source of rules which should be applied in future cases.

Cases can convey information of several different types. Some, called *framework precedents* in [29], set out further tests for the application of a concept. Other precedents identify the features of a case which need to be considered when applying these tests. US Trade Secret misappropriation has, since its use in HYPO [2], been the most widely explored domain in the AI and Law literature [24], [3]. US Trade Secrets Law can be found in the *Restatement of Torts*, a treatise issued by the American Law Institute[3] which summarises the general principles of the common law governing torts in the United States. The relevant section is 757, *Liability for disclosure or use of another's Trade Secret* begins by setting out the general framework:

> "One who discloses or uses another's trade secret, without a privilege to do so, is liable to the other if
>
> (a) he discovered the secret by improper means, or
>
> (b) his disclosure or use constitutes a breach of confidence reposed in him by the other in disclosing the secret to him."

It then goes on to state what must be considered to apply these principles, for example to determine whether information should be considered a trade secret:

> "Some factors to be considered in determining whether given information is one's trade secret are: (1) the extent to which the information is known outside of his business; (2) the extent to which it is known by employees and others involved in his business; (3) the extent of measures taken by him to guard the secrecy of the information; (4) the value of the information to him and to his competitors; (5) the amount of effort or money expended by him in developing

[3] https://www.ali.org/publications/show/torts/

> the information; (6) the ease or difficulty with which the information could be properly acquired or duplicated by others."

As well as identifying the aspects that need to be considered, precedents will also discuss the significance to be accorded them in various circumstances [30]. Thus if the plaintiff had disclosed the information to outsiders, the precedent would consider whether the *extent* of the disclosures gave a reason to find for the defendant. Similarly, the *Restatement of Torts* quoted above identifies as one factor to be considered "the ease or difficulty with which the information could be properly acquired or duplicated by others." Whether or not this factor applies in a particular case can be the subject of dispute, and some decisions suggest how such disputes may be resolved. An example of an argument at this level can be found in *Technicon Data Systems Corp. v. Curtis 1000, Inc*[4]: "The Court reasoned that the process had required over two-thousand hours, and still had not yielded a fully functional product. The Court held that this amount of time indicated that a trade secret was not readily ascertainable." This suggests that the time taken to reproduce the information is an important consideration, and the suggested threshold should be respected in future cases when determining whether this factor is present.

Note that these are factors which need to be taken into account and, since there will typically be factors for both sides, weighed against each other: they cannot be interpreted as sufficient conditions. This gives rise to a third role for precedents, the one which has received the most attention in AI and Law, starting with the CATO syatem [31]. Where there are factors for both sides, precedents establish preferences between sets of factors. Thus we may find in a decision a ruling which determines the appropriate outcome in a case in which several of the above factors are present. An example with factors (3) and (6) in the quotation above is *Mason v. Jack Daniel Distillery*[5]:

[4]Technicon Data Systems Corp. v. Curtis 1000, Inc., 224 U.S.P.Q. 286 (Court of Chancery Delaware, New Castle County 1984).

[5]Mason v. Jack Daniel Distillery, 518 So.2d 130 (Ala.Civ.App.1987).

> We note that absolute secrecy is not required ... "a substantial element of secrecy is all that is necessary to provide trade secret protection." Drill Parts, 439 So.2d at 49. In this regard, we note that courts have protected information as a trade secret despite evidence that such information could be easily duplicated by others competent in the given field.

This expresses a preference for the plaintiff's security measures over the possibility of reverse engineering. That the preferences between sets of factors found in cases was can be expressed as a set of rules was shown in [32], and formal models of this aspect of precedential constraint have been proposed in [33] and [34].

2.1.3. Special Requirements on Legal Knowledge

There are two particular aspects of the legal domain that need to be given particular consideration when representing legal knowledge: explanation and ease of maintenance.

Explanations are crucial in legal systems [17]. When presenting an argument in court, a simple assertion that one's client should win is useless: one must present the reasons why one's client should win. Thus for intending litigants, it is the explanation that will enable them to present their case. Further, the parties to a case have a right to explanation when the case is decided [16]. The loser of a suit has a right to know why they lost, and if they are not satisfied with the explanation there is, except at the highest level of Court, a right to appeal. Explanation is necessary if justice is not only to be done but to be seen to be done. Explanations are often based on the Issue-Rule-Application (IRAC) method of legal analysis, IRAC, or variants on it, is widely taught in law schools[6]. The key point about IRAC, and its variants, is the notion of

[6]For one example, se City University of New York (https://www.law.cuny.edu/legal-writing/students/irac-crracc/irac-crracc-1/). Use of IRAC is advocated by the LexisNexis survival guide for law students available at https://www.lexisnexis.co.uk/students/law/.

issue: typically it is one particular point in a case that is in dispute. How this is resolved is what needs explanation: the other aspects of the case which are accepted by both parties and are not in dispute need no discussion. IRAC had been advocated for use in AI and Law systems in [35].

The second important feature of legal knowledge is that it changes. If one is building a medical system, one can do so with confidence that the human body is not going to change (although, of course, our understanding of it may increase). In contrast, laws are in a constant state of revision and while some revisions may be small, others may be quite dramatic. Moreover, we find that case law also tends to change over time. Decisions in legal cases are supposed to reflect social attitudes and as attitudes change we find that emphasis may be placed on different considerations[7]. Also there may be a *landmark* decision which introduces a new consideration or overturns an established principle, and requires a reinterpretation of the existing understanding of case law. An example given in [36] is the the case of *Carrol v US*, which introduced the "automobile exception" to the US 4th Amendment. Such changes present a particular problem for machine learning systems (see [5], which reports an experiment showing how using older data in the training set degrades performance), but also means that conventional systems must constantly reflect such changes in their representation, making ease of maintenance of crucial importance.

2.1.4. Layers of Legal Reasoning

The structure of legal knowledge as described above, indicates that there are a number of layers of legal reasoning: a number of steps that must be gone through to move from the evidence presented in a case to a decision. The role of intermediate predicates, predicates that represent legal concepts that mediate between facts and legal consequences, has long been recognised in both the jurisprudence and AI and Law literature [37], [38] and [39]. In [39], factors are

[7] See Justice Marshall's remark in his opinion of *Furman v Georgia* that "stare decisis must bow to changing values".

seen as playing the role of these intermediate concepts. It is, however, possible to take a finer grained view as in [40] and [41]. There the reasoning starts by moving from the *evidence* presented to the *facts* as accepted by the court. On the basis of these facts, *factors* are ascribed. The "balance of factors" [22] can then suggest how the various *issues* pertinent to the legal question under dispute should be resolved. Once the issues have been resolved, the *outcome* of the case follows from a logical model of issues [42], found in statute or the relevant framework precedents. There is, therefore, a sequence of steps that must be gone through when considering a legal case. In AI and Law, different systems have addressed different parts of this sequence. The parts addressed by some leading systems in AI and Law is shown in Table 1.

Table 1: Layers of Statements in a Legal Decision and Some Example Systems

	BNA [27]	HYPO [2]	CATO [24]	IBP [42]	Bex [43]	NIHL [44]
Outcome	X			X		X
Issues	X		X	X		X
Factors		X	X	X		X
Facts		X			X	X
Evidence					X	

2.2. Approaches to Representing Legal Knowledge

For a detailed account of how various approaches to knowledge representation used in AI and Law support explanation see [17]. The main approaches found in the AI and Law literature are

- *Rule based approaches* (e.g the British Nationality Act program of [27]). Given the definitional nature of statute law, the rule based paradigm presents a natural choice for representing such knowledge. It is, however, less suitable for the lower layers of legal knowledge and systems using this paradigm typically assume that the users will be able to supply the required knowledge of case law.

- *Factor based approaches* (e.g. HYPO [2] and CATO [24]). These approaches offer a direct way of representing the intermediate concepts which emerge from case law as described in Section 2.1.2. However, they do not take full advantage of the structure provided by the statutes. Thus in formal accounts of this approach such as [33], irrelevant distinctions may unduly affect the reasoning.

- *Hybrid approaches* using both rules and factors (e.g. CABARET [45] and IBP [42]). These systems use rules at the top level and then interpret the undefined terms in the rules using factor based reasoning. This enables the domain structure to be exploited.

- *Argumentation approaches* (e.g. [46]) models the reasoning of CATO as a repertoire of argumentation schemes, and so also covers factors and issues. An argumentation approach is also used in [43] to move from evidence to accepted facts. These approaches support a very natural form of explanation using terminology familiar to users.

- *Machine Learning approaches* (e.g. [4] and [5]). These approaches do not use any representation of the law, but build a predictive model based on large numbers of previously decided cases. One major deficiency of current approaches is that they are unable to give a justification of their reasoning in terms of appropriate legal concepts[8].

From the various approaches, a number of desiderata for a representation of legal knowledge emerged:

- A clear need to respect the hierarchical nature of legal knowledge. This

[8]The need to explain reasoning from these systems has led to interest in so called Explainable AI (XAI), e.g. [47] and [48], and Argument Based Machine Learning [49]. In AI and Law, explanations of machine learning systems have attempted to draw on established symbolic techniques, either by learning to ascribe factors [14], [50], or by providing an independently generated explanation [51] [15]. Note that these approaches need to supplement the learned model with a symbolic model of the domain.

is so that the layers of different types of knowledge shown in Table 1 can be kept separate but appropriately related. The abstract factor hierarchy introduced in [24] is a good example of what is needed.

- A second role for the hierarchy is to split the overall question into a series of issues. This was shown to be of importance in hybrid systems such as [42] and [52], and recently emphasised in [30] and [53].

- It is important to be able to represent different styles of reasoning. As revealed in hybrid systems such as [42] and [26], sometimes rule based reasoning with necessary and sufficient conditions will be appropriate, but at other times balance of factors and purposive reasoning may be required. Moreover, if we allow non-boolean factors, as in [54] and [55], we may need additional techniques to allow more arithmetical reasoning [56] [57].

- The representation must be capable of adapting to change, especially to changes driven by evolving case law. There will also be changes consequent on legislative amendment. The important thing is to be able to identify, and keep, those parts of the representation unaffected by the changes. The key to a maintainable representation is modularity, so that any changes to the law can be associated with specific parts of the representation, and any changes to the legislation localised to a particular module [58].

- The representation must support effective explanations. Argumentation based explanations, both those based on precedent cases such as [24] and those based on argumentation schemes such as [46], have been able to provide effective explanations and they have also been adopted to provide explanations for machine learning, as in [15].

The ANGELIC methodology [19], [44] was developed to fulfill these requirements. It forms the basis of our approach to modelling the ECHR domain, which is the case study used in this paper, and will be described in Section 2.3. For a detailed account of how the ANGELIC methodology used in this paper developed from previous work, see [3].

2.3. The ANGELIC Methodology for Legal Reasoning

The central idea of the ANGELIC methodology is to base the representation on the Abstract Dialectical Frameworks (ADFs) of Brewka and Woltran ([18] and [59]). Although originally restricted to three valued nodes (true, false and undecided), ADFs were further generalised to *weighted* ADFs in [60], to accommodate real numbered values between 0 and 1 for the nodes. In Dung's AFs [61] nodes are linked by an *attack* relation and a node is acceptable if and only if none of its attackers (which we will call its *children*) are acceptable. ADFs generalise this so that while the status of a node is still determined by the status of its children, this is done using acceptance conditions local to the node. The definition of a weighted ADF in [60] is:

Definition 1. A weighted ADF (wADF) over V is a tuple $D = (S, L, C, V, \leq_i)$, where

- S is a set (of nodes, statements, arguments; anything one might accept or not),
- $L \subseteq S \times S$ is a set of links,
- V is a set of truth values,
- $C = \{C_s\}_{s \in S}$ is a collection of acceptance conditions over V, that is, functions $C_s : (\text{children}(S) \to V) \to V$,
- $(V_u\ , \leq_i)$ – where $V_u = V \cup \mathbf{u}$ – forms a complete partial order with least element u.

For legal purposes we can specialise this definition. We choose *statements* from the options as to what nodes represent, capturing that a party is favoured by the outcome, that an issue is resolved, that a factor is present, that a fact is accepted, etc. We also restrict ourselves to real numbers in range 0 ...1 as truth values, and so can dispense with the final clause of the definition.

This structure proves ideally suited to representing the legal knowledge described in Section 2.1. The factor hierarchy of CATO [24] conforms to this

structure: the issues, abstract factors and base level factors are all statements, and the status of non-leaf nodes (base level factors are givens) is determined exclusively by their children. We now, however, also have the ability to associate acceptance conditions with each node. These conditions are specified for each node, and so can allow for the acceptance of different nodes to be determined differently.

The ability to specify acceptance conditions appropriate to each node means that we can specify them as necessary and sufficient conditions or prioritised sufficient conditions derived from precedents with a default to enforce burden of proof[9], depending on whether rule based reasoning or balance of factors reasoning is appropriate, using either of the result or the reason model presented in [33]. The flexibility afforded by acceptance conditions particular to specific nodes becomes even more useful when we allow non boolean nodes, to represent the extents and amounts required for factors such as those mentioned in the *Restatement of Torts* as quoted in section 2.1.2 above. As well as functions such as *maximum* and *minimum* that enable fuzzy disjunction and conjunction [62], other functions such as comparison with thresholds, weighted sums, and equations representing trade-offs [57] have been used. A reconstruction of CATO with non-boolean factors is described in [63].

The top three layers from Table 1, which correspond to the abstract factor hierarchy of [24], can be represented using these techniques. The ANGELIC methodology, however, extends this hierarchy with an additional layer to represent the facts on the basis of which the factors are ascribed. The facts are intended to be obtained from the user, and so the leaf nodes of the ADF are *questions* designed to elicit the relevant facts. These nodes are associated with a textual question to put to the user. The user answers with "true", "false", or a number between 0 and 1, and the node assumes the value supplied. This means that the final layer, evidence, is not represented: the user is expected to

[9]In law, for each proposition, the party that is required to show it is said to have the 'burden of proof'. If that party is unable to show it, the proposition is deemed to be false.

assess the evidence and supply the accepted facts. On the basis of this, factors can be assigned, issues resolved and the outcome determined.

Note that this approach does require the user to supply the facts. But, of course, even if the system was to apply machine learning techniques to natural language, *someone* will have needed to draft the natural language description of the case. If the input comes from the facts section of a judgement as in [4], the description used will have been drafted by the trial judge. If using some pre-trial statement of facts as suggested in [64], the facts will have been drafted by some legally qualified employee of the court. Thus our questions impose no greater demands than these systems: indeed the questions provide a structure which supports the task of describing the case. Thus if incorporated in a setting where the expertise is available, for example supporting the trial judge, answering the questions imposes no additional resource requirements than does drafting the statement of facts used by the machine learning approaches.

The ADF supplies the desired modularity, since each node is determined exclusively by the status of its children. Thus rules cannot conflict: only the rules within a node are active at any given time and the conflicts between rules within a node are resolved by their priority ordering. A new factor can be included by adding a child, and a new precedent by adding the rules, or changing the priorities between existing rules, required to express the decision in that case, with full confidence that there will be no unwanted consequences elsewhere in the hierarchy. For a discussion of maintenance issues, see [65]. A detailed example of implementing a change in response to an unexpected decision in a case is given in [66].

Turning to explanation can take the form of an argument. Each parent is the conclusion of an argument with its children as precedents. The children can in turn each be established by an argument with their children as premises. We can thus construct a series of subarguments, until we reach the leaf nodes, where the answers given by the user are accepted without further argument. This argument-subargument structure, bottoming out in accepted facts, corresponds to the structured arguments of ASPIC+ [67], as described in [68]. This enables

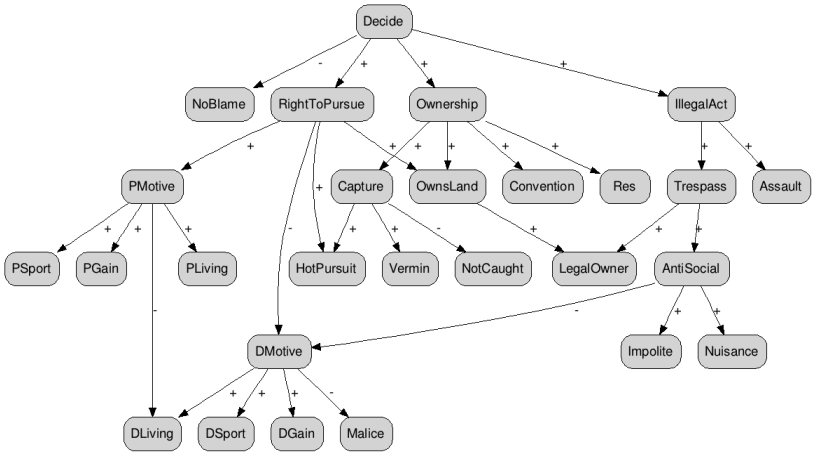

Figure 1: Visual representation of the ADF for the Wild Animals domain, as described in detail in [19].

the explanation to be produced as an argument from the facts to the case outcome [69].

The *ANGELIC* methodology has been used to model a variety of domains. As regards academic domains, the US Trade Secrets domain of HYPO and CATO, the much smaller wild animals property law domain introduced in [25] and the widely discussed automobile exception to the US 4th Amendment (e.g. [70]) were all represented in [19]. For illustration purposes, Figure 1 shows a visual representation of the ADF constructed for the wild animals domain.

In subsequent work, the CATO domain was remodelled with non booleans in [71] and the methodology has also been applied in a commercial environment to a variety of domains in collaboration with the large law firm, Weightmans. The most notable of the projects with Weightmans related to cases regarding Noise Induced Hearing Loss claimed to be due to employer negligence [44][10]

Most recently the ANGELIC methodology has been applied to the ECHR, and the work on this domain will be the subject of the remainder of this paper.

[10]Weightmans received the 2019 Eclipse Modern Law Award for best use of technology for this work.

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

The European Convention on Human Rights (ECHR) is a regional human rights treaty that is now ratified by 47 European states[11] and covers almost the whole of Europe. The most important feature of this treaty is the so-called right to an individual application. In other words, all alleged victims of human rights violations[12] can bring an application to the European Court of Human Rights (ECtHR) against the Member State that allegedly violated their rights. The ECtHR then can declare that the respondent State has violated human rights. In this case, the ECtHR delivers a judgment finding a violation of the Convention and the Respondent state has to pay monetary compensation to the applicant and change its law and practice to prevent similar violations from happening in the future.

Since 1960, when it was established, the ECtHR has delivered judgements in thousands of cases and created a significant body of legal precedents. Although the Court's previous judgments are not binding on the Court, it does not depart from them without a very good reason and does so only in a very small number of cases. Therefore, it is safe to try to predict the outcomes of the pending applications by referring to the Court's previous case law. The ECHR enshrines rights that can be divided along different lines. Some of the rights are absolute, which can never be interfered with. Prohibition of torture (Article 3 ECHR) is one among such rights. Any *interference* (the term used in the ECHR) of this right will lead to a violation. Some other rights are more complicated and some interference is possible but such interference needs to comply with strict rules established by the ECHR and ECtHR. For example, Article 10 enshrines freedom of expression. Some expression that can cause harm can be prohibited

[11] Effectively, the only European country which has not ratified the European Convention on Human Rights is Belarus.

[12] A victim may be a child, an adult, a group of people, a church, a company or a political party.

but such prohibition cannot be more than is necessary in a democratic society. Right to a fair trial (Article 6 ECHR) is not absolute: some limitations are possible in certain circumstances but overall the member states are required to ensure that the parties to a civil claim or the defendant in a criminal trial are treated fairly. In the sense of Article 6, *fairness* has a specific meaning. It does not mean that the outcome of the case must be universally accepted as fair – it is difficult to measure what fair may mean to different parties. Fairness here has a much more formal meaning. In this sense, it means that the case should be dealt with by an independent tribunal on the national level, that the parties to the case have equal rights, that those accused know what they are accused of, that they also have access to legal aid and if they do not understand the language of the process then an interpreter should be provided. If any of these entitlements are not provided, then the ECtHR can find a violation of Article 6. Thus the concern is for procedural fairness, rather than distributive fairness.

Article 6 is the most used Article of the ECHR; the majority of the applications submitted to the ECtHR complain about a violation of Article 6. Considering that the Court's backlog is one of the key challenges that the ECtHR is facing now, more automation enabling speedier resolution of the applications is of crucial importance. Although the number of pending applications reduced since 2010 when it reached 150,000, it was still over 60,000 in 2021 [72]. It has been estimated that the Court will need years to sort out its backlog even if the influx of new applications were to stop.

The ECHR has proved very popular for experimentation with machine learning techniques for legal judgment predication tasks; for example, see [4], [5], [6], [8] and [7]. These studies all report success, with correct predictions being achieved in around 70-85% of cases, which is arguably unacceptably low for practical use. JURI Says, the program described in [8], reports a success rate of 55.9% over the last year, although it reached 76.9% for February 2022)[13].

[13]JURI Says can be found at https://jurisays.com/ (accessed 2022/07/26).

4. ADF Model Design – Legal Foundations

We first developed an ADF model, extending that produced in [69], covering the whole of Article 6. We then, however, focussed specifically on whether an application to the ECtHR is admissible or not, which is itself a substantial task. All applications submitted to the ECtHR need to be admissible in order to be considered on merits. In other words, the Court needs to establish that the application complies with a set of formal rules before it can examine the substance of this application [73]. The set of these rules is enshrined in Articles 34 and 35 of the ECHR. These rules were elaborated in the *Practical Guide on Admissibility Criteria* prepared by the Court [74]. The *Practical Guide* was used to inform the current model.

Although the process of considering admissibility of applications is often presented as a binary choice between admissibility and inadmissibility that does not require any judicial discretion, this view is not completely adequate. The process of admissibility still requires some assessment of law and facts and in some cases, judicial discretion [75]. Having said that, determination of admissibility is a much more formal process and it is much easier to describe in precise terms than consideration of merits.

Admissibility includes two types of rules: first, the ECtHR needs to establish that the application falls within its jurisdiction. In other words, the Court needs to confirm that it can deal with this application. For example, it needs to be established that the applicant brought an application against one of the member states, the alleged violation of human rights took place after the ECHR was ratified by the respondent state, that the application has been submitted by the victim of a violation or their relatives and that the application is only concerned with the rights that are enshrined in the ECHR. If any of these conditions is not satisfied, the Court will have to declare the application inadmissible.

Secondly, the ECHR established a set of formal rules that the application itself needs to comply with. These rules for example, include that the application was first submitted at the national level and was rejected by the national judicial

bodies and that it should be submitted within 6 months after the highest judicial body rejected the same application on the national level. This application should not be abusive, anonymous or trivial. This application also should not be clearly without merits or – in the ECHR terms – manifestly ill-founded. Again, if these conditions are not satisfied, the Court declares an application inadmissible. The Court's decision as to inadmissibility is final and cannot be appealed against.

The importance of admissibility is often underestimated. On average about 90% of all applications submitted to the ECtHR are declared inadmissible. This rate varies across states and years but in all cases, it is very high. For instance, in 2019, 44,500 applications were submitted to the Court and in the same year 38,480 applications were declared inadmissible. At the same time, in 2019, the Court delivered only 2,187 meritorious judgments [76]. This statistic is illustrative only: some of the submitted 44,500 applications were not dealt with in 2019 and joined the Court's backlog: most of the judgments delivered in 2019 were dealing with the applications that were submitted before 2019. On average, however, these numbers are quite telling: a major number of applications is declared inadmissible every year, so our project has potential importance for both the applicants who might want to avoid inadmissibility and for the Court for which consideration of inadmissible applications takes a significant proportion of its time and resources which could be re-allocated to the meritorious cases and so reduce the backlog. In the next section, we describe the implemented tool that we have produced to enable decision support for the important issue of admissibility of cases submitted to the ECtHR. The model used in the tool captures the factors discussed above that need to be examined to determine admissibility and is a result of close consultation with our expert on the ECtHR.

5. ADF Implementations

In this section we describe how the ADF model is implemented. There are two implementations that we describe; firstly an implementation that handles predictions of Article 6 cases implemented in Prolog, secondly an implemen-

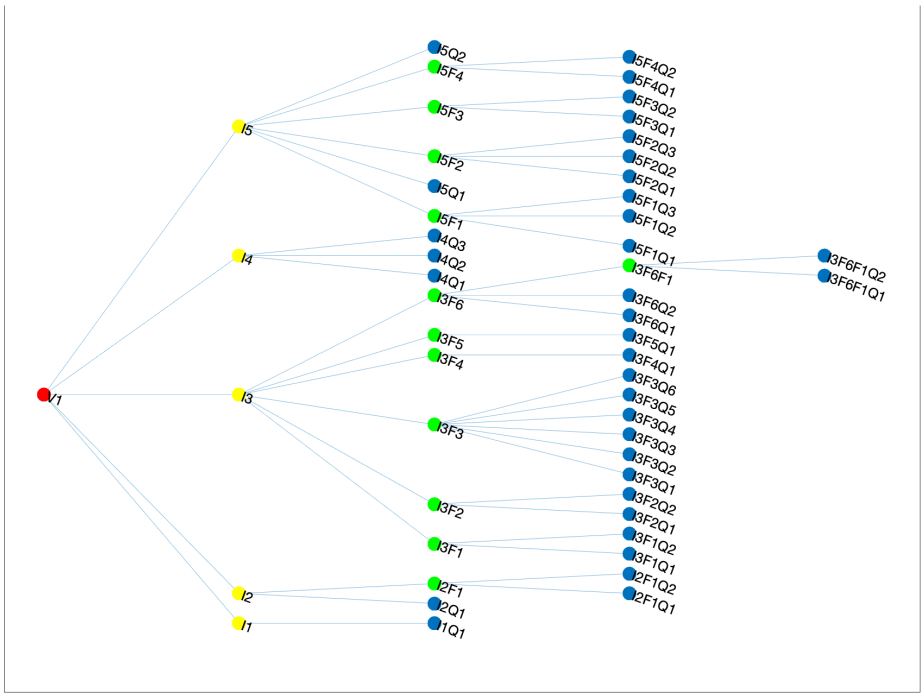

Figure 2: Full tree representation of the ADF used in the prototype for Article 6, given in more granular tabular form in Appendix A.1. The verdict (the program's recommendation) is shown in red, the issues are shown in yellow, abstract factors are shown in green, and base level factors are in blue. The labels signify the ID of the node.

tation that handles admissibility of European Court cases implemented as a website.

5.1. Article 6 Implementation

To develop the ADF, we researched the lawyers' guide to Article 6 [77, 78][14]. We followed the ANGELIC methodology described in Section 2.3, and so the ADF represents a hierarchy containing the various elements identified in Table 1. The main issues which are discussed when finding a violation of Article 6 are

[14]We do not distinguish between the criminal and civil limbs, consistent with the previous literature.

identified. Each issue then has a number of factors that describe whether the issue is relevant to the case in hand. Factors in turn can have sub factors that are identified. Once a factor can be answered with a question soliciting a particular fact, the path is ended. The question that is posed should be such as to be able to be answered by a lawyer familiar with the ECHR and the case under consideration, with a high degree of confidence. A visual representation of the ADF is shown in Figure 2 – see Appendix A.1 for full ADF details.

The ADF was then implemented as a Prolog program, which can be run on the command line. The Prolog code traverses each node in the ADF, firstly evaluating if the node is accepted or rejected, then printing a human readable explanation.

The following is a sample code snippet for the node *isFairAndPublic* (issue I3 in Appendix A.1), which concerns whether the case was fair and public.

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

When traversing the nodes of the ADF, the Prolog code will need to evaluate the leaf nodes. These represent the base level factors and can be instantiated by asking the user the corresponding question. The following is the code snippet for the leaf node *notPreventedFromLawyerAccess* (I5F2Q3 in Appendix A.1), representing whether the defendant was prevented from accessing lawyers. For

example, suppose the question is "Was the defendant prevented from accessing a lawyer?"; if this true, then the base level factor f40 is added to the list of factors present in the case.

```
notPreventedFromLawyerAccess(case(_,L), valid) :-
    member(f40, L),
    write("Not prevented from accessing lawyers"), nl.
notPreventedFromLawyerAccess(case(_,L), invalid) :-
    not(member(f40, L)),
    write("Prevented from accessing lawyers"), nl.
```

The user will need to answer all questions in order to provide the full list of base level factors to the program. An example of a full case input by the user is:

```
violationOfArticle6(case('MARGUÅ v. CROATIA JUDGMENT',[
    f2,f4,f6,f7,f9,f10,f12,f13,f15,f16,f17,f18,f19,f20,f21
    ,f22,f25,f26,f27,f28,f32,f33,f35,f36,f37,f38,f39,f40,
    f41,f42])).
```

When the Prolog code has finished executing, each node has been evaluated with the resulting output making up the explanation and outcome of the case. Example output for a case showing no violation is given in Appendix A.3.

The output shows the final result on the last line, which in this case is that there is no violation. The reasoning as to how this decision was reached is given in the proceeding lines, and can be read from top to bottom, with the different issues indicated by indentation.

5.2. Admissibility Implementation

The ultimate aim of the admissibility program is to provide the public, assisted by non-specialist lawyers, with the ability to get a recommendation on whether the case they want to submit to the ECtHR would be accepted as admissible. The ADF designed for this task would therefore need to expand upon the previous implementation to give assistance in answering the high level

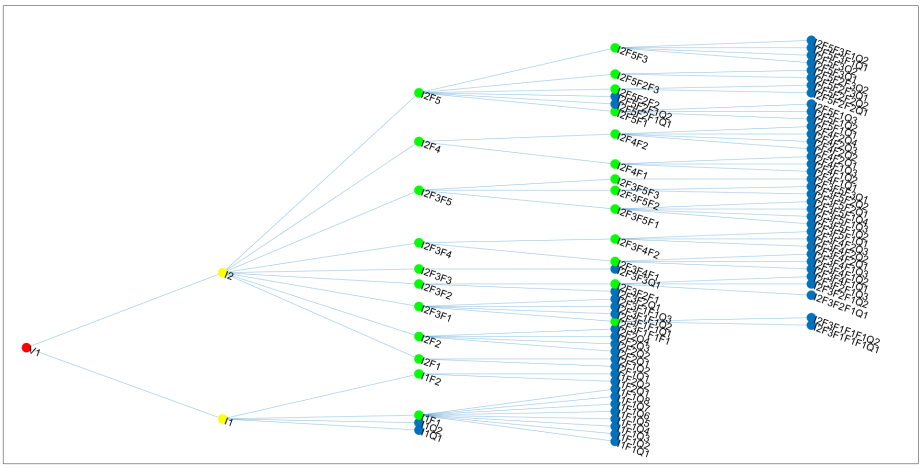

Figure 3: Full tree representation of the ADF for Admissibility used in the prototype. The verdict (the program's recommendation) is shown in red, the issues are shown in yellow, abstract factors are shown in green, and base level factors are in blue. The labels signify the ID of the node.

questions of that implementation. Taking the same approach as before, we consulted the lawyers' guide to admissibility [74] to gather the issues, factors, and questions we needed to create the ADF. A visual representation of the ADF for Admissibility is shown in Figure 3, and the full ADF is given in Appendix A.2.

To assist the end users, who are envisaged to be non-computer science experts, in using the implementation, we have created a GUI-based tool to enable easy use. The front end of the program was implemented both as a JAVA program and as web based tool. The code implements the ADF, with the web-based tool that poses questions to users being implemented in JavaScript. How the questions are posed to the user can be seen in the screenshot shown in Figure 4. The order of the questions was determined by input from our legal expert and the accepting logic of the factors.

The user continues to answer questions until the ADF can be resolved. If the ADF recommends to submit the case then a full explanation is given, as in

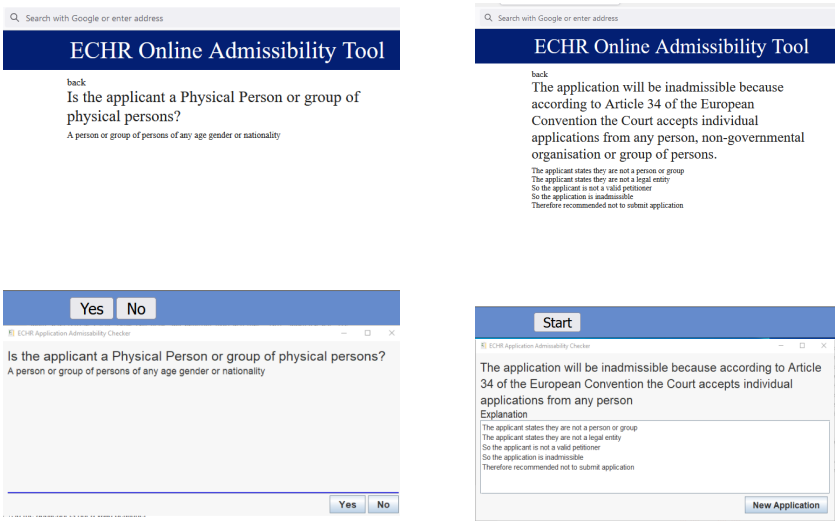

Figure 4: Admissibility program showing an example question to the user (Left), and showing a recommendation to not submit the case (Right). The program is shown as a web based tool (top) and as a JAVA program (bottom).

the previous Article 6 Prolog implementation. When there is a recommendation to not submit the case, the program presents the reasoning for how it came to the decision and why the recommendation not to submit is given. However in this case the full reasoning is not given, as it is plain to see from the Article 6 Prolog program that the amount of information is more than can easily be absorbed by a lay user. In order to ensure that the intended user, who is not a lawyer or computer expert, can parse the information, only the relevant part of the explanation is shown. The results of this can be seen in the screenshot in Figure 4.

Short explanations were generated by taking the last question that was asked, which made the application inadmissible. From the parent of the question node, all children that have had their associated question answered will be part of the explanation. We then traverse back up through the tree from the parent back to the root node. Each node prints its status in human readable form. Thus generating our explanation.

Consider an example for generating the shorter explanation using the Ad-

missibility ADF where the program has presented the user with the question "Is the applicant a Physical Person or group of physical persons?" (I2F1Q1), to which the user has provided the response "no". The program now asks "Is the applicant a legal entity" (I2F1Q2), to which again the user provides the response "no". The program can now resolve that the applicant is not a valid petitioner (I2F1), and that the application is inadmissible (I2), and therefore the program recommends not to submit the application (V1). As the program traverses the ADF back to the root node, each node prints a sentence which presents the information in a human readable form. The explanation generated is shown in Figure 4.

6. Validation and Evaluation

6.1. Overview of evaluation activities

Our evaluation activities cover three different aspects. Firstly, we determined the accuracy of the Article 6 Prolog model, by evaluating a total of 40 cases in our model and examining whether the program produces the correct output.

The wider aim of our work is to bring AI tools to the law community that practitioners themselves would find useful. This gives our motivation for our second evaluation exercise; specifically, determining the admissibility of a case is an aspect of the ECtHR which shows promise as a practical application of such such tools since the task of determining admissibility is carried out for every case submitted and is a major factor in the current large backlog of cases, as noted in Section 3. To determine whether our admissibility program is useful and appropriate, a pilot study was conducted to collect feedback from a select group of lawyers who were given access to the JAVA program.

Finally, a wider usability study was conducted where members of the law community were asked to evaluate the web version of the admissibility program. Below we report outcomes from all three evaluation exercises. For the two pilot studies that involved human participants, we made two formal ethics

applications to, and were subsequently granted approval from, the University of Liverpool's Research Ethics Committee[15].

6.2. Validating the Accuracy of the Model

Our first requirement is to determine the accuracy of the ADF model we have developed; accordingly, this sub-section describes the experimental evaluation of our Prolog implementation of Article 6.

For this exercise, we first validated the results of our Prolog implementation using a set of 10 cases that were used to evaluate an earlier version of our ADF model reported in [69] that did not cover the determination of admissibility in depth, as we have done in the implementation detailed in this paper. We then conducted our main evaluation activity using a new set of 30 Article 6 cases whose judgements are released through HUDOC[16]. In total we have manually transcribed the factors from the 30 Article 6 cases, with a breakdown of 15 cases where there was found to be an Article 6 violation and 15 non-violation cases. This set of cases are in the corpus used in Aletras *et al.'s* work [4]. Our full suite of test cases is given in Appendix A.4.

The results for the violation cases are shown in Table A.5 of Appendix A.4, in which 14 out of the 15 cases are correctly predicted as violations. The results for the no violation cases are shown in Table A.6 of Appendix A.4, in which all 15 cases are correctly predicted as non-violations; our model has no violation as a default, so there were no factors present in any of the cases that would make the model predict a violation.

The results of the experimental evaluation on the 30 new cases show that the Prolog program gave an incorrect result on only one of the cases in the corpus.

Analysing where the issues arose in that one incorrect case, we note that the case is annotated as a level 1 case. Level 1 cases are cases that "significantly contributes to the development of case law"[17]. The case therefore is considered

[15] Project references: 8536 and 8737

[16] HUDOC can be found at `https://hudoc.echr.coe.int/`

[17] `https://www.echr.coe.int/Documents/HUDOC_FAQ_ENG.pdf`

to be non-trivial by the lawyers and judges who ruled on that case, which adds to the difficulty in ascribing the factors to use in the ADF.

The case in question is Karanović v. Bosnia and Herzegovina[18]. We note that the decision of the program relied on a single factor, which if invoked would have produced a correct result. In this case f10 was incorrectly ascribed, and should not have been included. Thus, the incorrect outcome can be attributed to failing to correctly ascribe one factor, highlighting the importance of reliable factor ascription when constructing the model. Nonetheless, our explainable model has greatly facilitated traceability of the reasoning to determine where and why any incorrect predictions are returned. As a further note, the ADF model that we have developed for Article 6 does not provide guidance on how questions should be answered. Further development of the ADF model, where base level factors are deepened to provide more guidance, would allow for the questions to rely less on expert judgement, improving the reliability of factor ascription.

Overall the accuracy of our model, evaluated on a total of 40 Article 6 cases, is 97%. These figures compare with the 79% accuracy over 584 cases reported for the Aletras *et al.* model [4], though that work covered Articles 3, 6, and 8 of the ECHR. If only the 80 Article 6 are compared, the accuracy result is 84%. Thus, our model shows high accuracy, which comes at the expense of speed of processing each individual case, since we have to ascribe factors for each new case. We have constructed our model from law and guidance documents that the ECtHR has provided, rather than analysing individual cases, as required for the case-based models such as [24] and [42]. By building the model in this way we have achieved high levels of accuracy without the high cost of resources needed when building models which require the analysis of many past cases. Our belief is that lawyers and judges will first want to see high accuracy which can be

[18] `https://hudoc.echr.coe.int/eng#{%22fulltext%22:[%22Karanovi%C4%87%20%20%20v.%20Bosnia%20and%20Herzegovina%22],%22documentcollectionid2%22:[%22GRANDCHAMBER%22,%22CHAMBER%22],%22itemid%22:[%22001-83372%22]}`

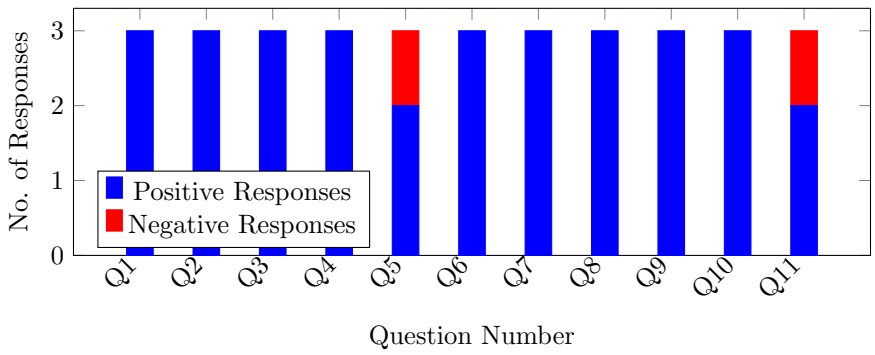

Figure 5: Graph showing number of feedback responses to each question in the pilot survey, as positive or negative feedback.

explained, then further development can speed up the processing of individual cases, rather than developing a model which aims for processing cases quickly from the outset, but which is incapable of explaining or justifying its reasoning.

6.3. Pilot Usability Study

The results reported in the previous sub-section show that our Article 6 model gives a high level of prediction accuracy, exceeding the accuracy provided by popular machine learning approaches. As we are aiming to produce tools that are useful to the law community, we now need to demonstrate that our implementations have a practical use. Thus we conducted a pilot study of the admissibility program described in Section 5.2. The pilot study used the JAVA version of the admissibility program, example screenshots of which can be seen in Figure 4.

The pilot study was conducted with a sample of our target audience, which is a small group of lawyers who work within the ECHR. The three lawyers who tested the prototype were asked to fill in a questionnaire that covers five different aspects of the prototype: functionality, usability, explainability, usefulness, and feedback on the questions used in the prototype tool. The set of questions and response options are shown in Appendix B.1.

The responses to the questionnaire have been condensed into positive or

negative responses, where the top two answers to a question are positive and the last two are negative. Figure 5 shows the number of positive and negative responses to each of the questions.

Though the results of the questionnaire come from a very small sample, with only three lawyers participating in the study, we were able to draw initial conclusions that the program developed worked well and was functional, since all the responses received on functionality (Q1, Q2) and usability (Q3, Q4) were positive. Another positive outcome is that two of the three ECHR lawyers responded that they found the justifications for the decisions sensible and understandable (Q5) and all three respondents agreed that the information was easy to parse (Q6). All respondents saw the usefulness of our prototype (Q7), with two respondents stating they would use the program as it currently stands, and the other affirming the usefulness but saying that some (as opposed to many) changes are needed. Again, all the respondents agreed that technology has a role to play in the legal domain (Q8): one respondent said technology is needed rapidly, while two cautioned that careful development will be needed.

The positive responses to Q9 and Q10 were particularly pleasing, since these questions directly concerned the central aims of this exercise: the users all agreed that the questions were suitable for them and that the program would save them time when assessing admissibility. While the majority of the feedback has been positive, it has also highlighted the need for domain experts to be a part of the development process (Q11): although two respondents felt the program reflected all or part of their own process of dealing with admissibility, one felt that only some aspects had been covered.

Overall the initial response to the program was very positive and indicated a sound basis for further dissemination and evaluation of our legal decision support tools. Encouraged by these results, we then extended the study to a larger group of potential users, further expanding our evaluation activities directly with the law community.

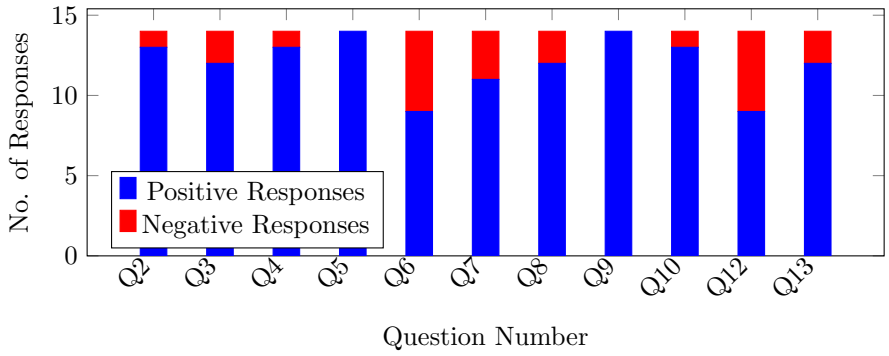

Figure 6: Graph showing number of feedback responses to each question for the second survey, as positive or negative feedback for the wider survey when using the admissibility program.

6.4. Evaluation with ECtHR Users

As the results of the pilot study were positive, we subsequently embarked upon a wider study to gather opinions from representative end users of our tool. We again sought participation from ECtHR lawyers, but for this exercise they evaluated the web based version of the admissibility program. The questions that were presented to the wider law community were the questions in the pilot survey plus a free text box for additional comments and an active consent question as part of our research ethics requirements. Appendix B.2 shows all the questions and possible response options.

A total of 14 lawyers completed the questionnaire, and each lawyer was able to claim a gifted £25 for completing the survey[19]. This recompense was not advertised to the lawyers and was only communicated to them when they had fully completed the survey. Figure 6 shows the results for the wider survey, showing the positive replies (first two possible response options) and the negative replies (last two possible response options). Questions 1 and 11 have been omitted as Question 1 is an active consent question and Question 11 is a free text box.

[19]The funds for this questionnaire came from a project supported by the University of Liverpool's Early Career Researchers and Returners Fund.

The wider responses are again very positive. The lawyers agree that the program runs well (Q2, Q3), and that the program is easy to use (Q4, Q5) and that the questions posed were easy to understand (Q10). Most of the lawyers agree that the explanations generated by the program justify the decision made (Q6), and the generated explanation was easy to read and understand (Q7). Encouragingly, almost all the lawyers found the application useful (Q8) and most agreed that the application would save them time (Q12). The respondents recognised the need for technology developments in the law domain (Q9) and found that our program reflects how they would process admissibility (Q13).

The wider community also provided feedback via a free text box, where they provided information on how they would like the program to be expanded. Some choice quotes from the lawyers include:

- "I think it is a good idea to develop legal tech tool for admissibility evaluation";
- "it would be more comfortable to have one checklist on one or two pages, rather than one-by-one questions";
- "It would be helpful if the program referred to the most important key cases and/or more substantive explanations from the jurisprudence of the ECHR regarding common reasons that justify inadmissibility."

From these comments we can see that this section of the law community is open using tech to help with their processing of admissibility decisions. Regarding our program specifically, they would like to see the program be further developed to be quicker to use and also for justifications to also include not only the literal explanation given but also why those explanations are correct by referring to jurisprudence of previous cases. This could be achieved without too much difficulty by associating each acceptance condition in the ADF with the relevant statutory clause or case from which it was derived.

Overall, we are highly encouraged by the results of our evaluation exercises as they show not only effectiveness of the reasoning models we have produced,

but also acceptance by stakeholders who expressed that the tools could be put to use in their work, given the transparency of the explanations provided.

7. Concluding Remarks and Future Work

The work described in this paper aims to support the legal community by providing AI-based tools to assist with improving access to justice. To achieve this, we have developed tools which present predictions of Article 6 cases and determine Admissibility of cases submitted to the European Court of Human Rights.

Using formal knowledge representation techniques, both tools presented have been designed as Abstract Dialectical Frameworks, modelled using the ANGELIC methodology [19]. The leading benefit of this approach is to ensure that the tools are able to explain why a prediction has been made. While individual cases are slower to process than equivalent machine learning approaches when building the model, there is an increase in accuracy of predicting cases. When there are changes to the law, our approach is easier to adapt to the changes than machine learning approaches, which will require retraining with minimal precedent cases.

The tool that predicts the outcomes of Article 6 cases showed very high levels of accuracy: the program achieved 97% accuracy over a total of 40 cases. Furthermore, we provided the admissibility program to three ECtHR lawyers to evaluate by completing a survey based on their experience using the program. The results of the initial small survey were promising, allowing us to expand the pilot study to a wider group.

From our interactions with lawyers who work in the domain of human rights cases, we focused our program on the task of determining the admissibility of cases. Deciding on whether a case is admissible to the court is a significant issue, highlighted by the large backlog of cases currently that await processing by the court. By refocusing our program on the admissibility issue, we were better able to meet our objective of providing tools that are appropriate and

useful to the lawyers and clients who are envisaged as users of such a system. This motivation is also behind our decision to transform our original Prolog tool into a web based system.

The web based tool that predicts the outcomes of admissibility cases was presented to a number of ECtHR lawyers. Encouragingly, they are open to such AI-based technology being implemented to help support their work. Using our tool they found that the admissibility program was easy to use and envisage that it would help save them time when processing admissibility cases, and in general that the program acceptably justifies the decisions it makes.

A strong basis has been given for further development of these tools by incorporating feedback from the lawyers into revisions. Future work will also be focussed on the development of new technical solutions to put machine learning approaches to use for the task of factor ascription (see [50] for initial steps in this line of work), with the ultimate aim of producing a hybrid system that reaps the benefits of machine learning for building the models and the benefits of knowledge representation techniques for reasoning over the models. Such a hybrid system would yield efficient decision support tools that meet the important criterion of providing much-needed explanations to the target end users. The work we have presented in this paper is a significant milestone on the path to this ultimate, long term aim.

Acknowledgements

We thank Kanstantsin Dzehtsiarou for his assistance with our modelling of the ECHR domain and for helping to organise the evaluation activities. We are also grateful to the lawyers working with the ECtHR who participated in our evaluation activities. The evaluation exercises were partially funded through the University of Liverpool's Early Career Researcher support fund, for which we express our gratitude. We also acknowledge partial funding from Research England through the Higher Education Innovation Funding (HEIF) Industrial Strategy fellows scheme.

References

[1] B. G. Buchanan, T. E. Headrick, Some speculation about AI and legal reasoning, Stanford Law Review 23 (1970) 40.

[2] E. L. Rissland, K. D. Ashley, A case-based system for Trade Secrets law, in: Proceedings of the 1st International Conference on Artificial Intelligence and Law, ACM, 1987, pp. 60–66.

[3] T. Bench-Capon, HYPO's legacy: introduction to the virtual special issue, Artificial Intelligence and Law 25 (2) (2017) 1–46.

[4] N. Aletras, D. Tsarapatsanis, D. Preoţiuc-Pietro, V. Lampos, Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective, PeerJ Computer Science 2 (2016) e93.

[5] M. Medvedeva, M. Vols, M. Wieling, Using machine learning to predict decisions of the European Court of Human Rights, AI and Law (2019) 1–30.

[6] I. Chalkidis, I. Androutsopoulos, N. Aletras, Neural legal judgment prediction in english, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 4317–4323.

[7] A. Kaur, B. Bozic, Convolutional neural network-based automatic prediction of judgments of the European Court of Human Rights, in: 27th AIAI Irish Conference on Artificial Intelligence and Cognitive Science, CEUR Workshop Proceedings 2563, 2019, pp. 458–469.

[8] M. Medvedeva, X. Xu, M. Vols, M. Wieling, Prediction system for the European Court of Human Rights, in: Legal Knowledge and Information Systems: JURIX 2020: The Thirty-third Annual Conference, Vol. 334, IOS Press, 2020, p. 277.

[9] T. Bench-Capon, The need for Good Old fashioned AI and Law, in: International Trends in Legal Informatics: A Festschrift for Erich Schweighofer, Editions Weblaw, Bern, 2020, pp. 23–36.

[10] F. Bex, H. Prakken, On the relevance of algorithmic decision predictors for judicial decision making, in: Conference for Artificial Intelligence and Law (ICAIL'21), 2021, pp. 175–179.

[11] C. Steging, S. Renooij, B. Verheij, et al., Discovering the rationale of decisions: experiments on aligning learning and reasoning, in: Proceedings of the Fourth International Workshop on Explainable and Responsible AI and Law (XAILA) 2021, CEUR Workshop Proceedings 3168, 2022.

[12] M. Medvedeva, M. Wieling, M. Vols, Rethinking the field of automatic prediction of court decisions, Artificial Intelligence and Law (2022) 1–18.

[13] K. Branting, B. Weiss, B. Brown, C. Pfeifer, A. Chakraborty, L. Ferro, M. Pfaff, A. Yeh, Semi-supervised methods for explainable legal prediction, in: Proc. of the 17th International Conf. on AI and Law, 2019, pp. 22–31.

[14] L. K. Branting, C. Pfeifer, B. Brown, L. Ferro, J. Aberdeen, B. Weiss, M. Pfaff, B. Liao, Scalable and explainable legal prediction, AI and Law (2020) 1–26.

[15] H. Prakken, R. Ratsma, A top-level model of case-based argumentation for explanation: Formalisation and experiments, Argument & Computation (2021)) 1–36.

[16] F. Doshi-Velez, M. Kortz, R. Budish, C. Bavitz, S. Gershman, D. O'Brien, S. Schieber, J. Waldo, D. Weinberger, A. Wood, Accountability of AI under the law: The role of explanation, arXiv preprint arXiv:1711.01134.

[17] K. Atkinson, T. Bench-Capon, D. Bollegala, Explanation in AI and Law: Past, present and future, Artificial Intelligence (2020) 103387.

[18] G. Brewka, S. Woltran, Abstract dialectical frameworks, in: Twelfth International Conference on the Principles of Knowledge Representation and Reasoning, 2010, pp. 102–111.

[19] L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, A methodology for designing systems to reason with legal cases using ADFs, AI and Law 24 (1) (2016) 1–49.

[20] D. N. MacCormick, R. S. Summers, A. L. Goodhart, Interpreting precedents: a comparative study, Routledge, 2016.

[21] K. D. Ashley, 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 Autonoma de México, Mexico City, 2004, pp. 1–30.

[22] N. Varsava, How to realize the value of stare decisis: Options for following precedent, Yale Journal of Law and the Humanities 30 (2018) 62–120.

[23] J. Hage, A theory of legal reasoning and a logic to match, Artificial Intelligence and Law 4 (3) (1996) 199–273.

[24] V. Aleven, Teaching case-based argumentation through a model and examples, Ph.d. thesis, University of Pittsburgh (1997).

[25] D. H. Berman, C. L. Hafner, Representing teleological structure in case-based legal reasoning: The missing link, in: ICAIL 1993, 1993, pp. 50–59.

[26] T. Bench-Capon, G. Sartor, A model of legal reasoning with cases incorporating theories and values, Artificial Intelligence 150 (1-2) (2003) 97–143.

[27] M. Sergot, F. Sadri, R. Kowalski, F. Kriwaczek, P. Hammond, H. Cory, The British Nationality Act as a logic program, Communications of the ACM 29 (5) (1986) 370–386.

[28] U. Schild, Open-textured law, expert systems and logic programming, Ph.d. thesis, Imperial College. London (1990).

[29] A. Rigoni, An improved factor based approach to precedential constraint, AI and Law 23 (2) (2015) 133–160.

[30] T. Bench-Capon, K. Atkinson, Precedential constraint: the role of issues, in: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law, 2021, pp. 12–21.

[31] V. Aleven, Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment, Artificial Intelligence 150 (1-2) (2003) 183–237.

[32] H. Prakken, G. Sartor, Modelling reasoning with precedents in a formal dialogue game, AI and Law 6 (3-4) (1998) 231–87.

[33] J. F. Horty, T. Bench-Capon, A factor-based definition of precedential constraint, AI and Law 20 (2) (2012) 181–214.

[34] H. Prakken, A formal analysis of some factor- and precedent-based accounts of precedential constraint, AI and Law 29 (4) (2021) 559–585.

[35] T. Bench-Capon, Explaining legal decisions using IRAC, in: Proceedings of Computational Models of Natural Argument (CMNA 2020), CEUR Workshop Proceedings 2669, 2020, pp. 74–83.

[36] E. L. Rissland, M. T. Friedman, Detecting change in legal concepts, in: Proceedings of ICAIL 1995, 1995, pp. 127–136.

[37] A. Ross, Tû-tû, Harvard Law Review (1957) 812–825.

[38] L. Lindahl, J. Odelstad, Intermediaries and intervenients in normative systems, Journal of Applied Logic 6 (2) (2008) 229–250.

[39] K. D. Ashley, S. Brüninghaus, A predictive role for intermediate legal concepts, in: Proceedings of JURIX, 2003, pp. 153–162.

[40] L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, Statement types in legal argument., in: JURIX 2016, 2016, pp. 3–12.

[41] J. Mumford, K. Atkinson, T. Bench-Capon, Machine learning and legal argument, in: Computational Models of Legal Argument 2021, CEUR Workshop Proceedings 2397, 2021, pp. 47–56.

[42] S. Brüninghaus, K. D. Ashley, Predicting outcomes of case based legal arguments, in: Proceedings of the 9th International conference on Artificial Intelligence and Law, ACM, 2003, pp. 233–242.

[43] F. Bex, H. Prakken, C. Reed, D. Walton, Towards a formal account of reasoning about evidence: argumentation schemes and generalisations, Artificial Intelligence and Law 11 (2) (2003) 125–165.

[44] L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, S. Whittle, R. Williams, C. Wolfenden, Noise induced hearing loss: Building an application using the angelic methodology, Argument & Computation 10 (1) (2019) 5–22.

[45] D. B. Skalak, E. L. Rissland, Arguments and cases: An inevitable intertwining, Artificial Intelligence and Law 1 (1) (1992) 3–44.

[46] H. Prakken, A. Wyner, T. Bench-Capon, K. Atkinson, A formalization of argumentation schemes for legal case-based reasoning in ASPIC+, Journal of Logic and Computation 25 (5) (2015) 1141–1166.

[47] A. Adadi, M. Berrada, Peeking inside the black-box: a survey on explainable artificial intelligence (xai), IEEE access 6 (2018) 52138–52160.

[48] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, et al., Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai, Information fusion 58 (2020) 82–115.

[49] O. Cocarascu, F. Toni, Argumentation for machine learning: A survey., in: COMMA, 2016, pp. 219–230.

[50] J. Mumford, K. Atkinson, T. J. M. Bench-Capon, Explaining factor ascription, in: S. E. (Ed.), Legal Knowledge and Information Systems - JURIX 2021: The Thirty-fourth Annual Conference, Vilnius, Lithuania, 8-10 December 2021, Vol. 346 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2021, pp. 191–196. doi:10.3233/FAIA210335.
URL https://doi.org/10.3233/FAIA210335

[51] J. Zeleznikow, A. Stranieri, The split-up system: integrating neural networks and rule-based reasoning in the legal domain, in: Proceedings of the 5th international conference on Artificial intelligence and law, 1995, pp. 185–194.

[52] M. Grabmair, Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs, in: Proceedings of the 16th International Conference on AI and Law, 2017, pp. 89–98.

[53] T. Bench-Capon, Using issues to explain legal decisions., in: Proceedings of the Fourth International Workshop on Explainable and Responsible AI and Law (XAILA) 2021, CEUR Workshop Proceedings 3168, 2022.

[54] J. Horty, Reasoning with dimensions and magnitudes, Artificial Intelligence and Law 27 (3) (2019) 309–345.

[55] A. Rigoni, Representing dimensions within the reason model of precedent, Artificial Intelligence and Law 26 (1) (2018) 1–22.

[56] M. Lauritsen, On balance, Artificial Intelligence and Law 23 (1) (2015) 23–42.

[57] T. J. Bench-Capon, K. Atkinson, Dimensions and values for legal cbr, in: Proceedings of JURIX 2017, 2017, pp. 27–32.

[58] F. Coenen, T. Bench-Capon, Maintenance of knowledge-based systems: theory, techniques and tools, Academic Press, 1993.

[59] G. Brewka, H. Strass, S. Ellmauthaler, J. P. Wallner, S. Woltran, Abstract dialectical frameworks revisited, in: Twenty-Third International Joint Conference on Artificial Intelligence, 2013, pp. 803–809.

[60] G. Brewka, H. Strass, J. Wallner, S. Woltran, Weighted abstract dialectical frameworks, in: AAAI 2018, 2018, pp. 1779–1786.

[61] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, Artificial Intelligence 77 (2) (1995) 321–357.

[62] L. A. Zadeh, Fuzzy logic and approximate reasoning, Synthese 30 (3) (1975) 407–428.

[63] T. J. Bench-Capon, K. Atkinson, Lessons from implementing factors with magnitude., in: Proceedings of JURIX 2018, 2018, pp. 11–20.

[64] M. Medvedeva, A. Üstun, X. Xu, M. Vols, M. Wieling, Automatic judgement forecasting for pending applications of the European Court of Human Rights, in: ASAIL & LegalAIIA 2021, CEUR Workshop Proceedings 2888, 2021, pp. 12–23.

[65] L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, Accommodating change, Artificial Intelligence and Law 24 (4) (2016) 409–427.

[66] T. Bench-Capon, T. F. Gordon, Implementing a theory of a legal domain, in: Legal Knowledge and Information Systems: JURIX 2016: The 35th Annual Conference, 2022, pp. 13–32.

[67] H. Prakken, An abstract framework for argumentation with structured arguments, Argument and Computation 1 (2) (2010) 93–124.

[68] K. Atkinson, T. Bench-Capon, Relating the angelic methodology and aspic+., in: COMMA, 2018, pp. 109–116.

[69] J. Collenette, K. Atkinson, T. Bench-Capon, An explainable approach to deducing outcomes in European Court of Human Rights cases using ADFs, in: Proceedings COMMA 2020, IOS Press, 2020, pp. 21–32.

[70] E. L. Rissland, Dimension-based analysis of hypotheticals from supreme court oral argument, in: Proceedings of the 2nd international conference on Artificial intelligence and law, 1989, pp. 111–120.

[71] T. Bench-Capon, K. Atkinson, Implementing factors with magnitude, in: Computational Models of Argument, IOS Press, 2018, pp. 449–450.

[72] European Court of Human Rights, Analysis of statistics, https://www.echr.coe.int/Documents/Stats_analysis_2021_ENG.pdf (2022).

[73] F. De Londras, K. Dzehtsiarou, Great Debates on the European Convention on Human Rights, Macmillan International Higher Education, 2018.

[74] European Court of Human Rights, Practical guide on admissibility criteria, https://www.echr.coe.int/Documents/Admissibility_guide_ENG.pdf (2022).

[75] A. Tickell, Dismantling the iron-cage: The discursive persistence and legal failure of a "bureaucratic rational" construction of the admissibility decision-making of the European Court of Human Rights, German Law Journal 12 (10) (2011) 1786–1812.

[76] European Court of Human Rights, Analysis of statistics, https://www.echr.coe.int/Documents/Stats_analysis_2019_ENG.pdf (2019).

[77] Guide on article 6 of the European Convention on Human Rights, right to a fair trial (criminal limb), https://www.echr.coe.int/documents/guide_art_6_criminal_eng.pdf (2021).

[78] Guide on article 6 of the European Convention on Human Rights, right to a fair trial (civil limb), https://www.echr.coe.int/documents/guide_art_6_civil_eng.pdf (2021).

Appendix A. ADFs

Appendix A.1. Article 6 ADF

ID	Prolog ID	Factor Text	Accepting Logic
V1	N/A	Article 6 Accepting	$I1 \wedge I2 \wedge I3 \wedge I4 \wedge I5$
I1	N/A	Is a victim	$I1Q2$
I1Q1	f2	Is the applicant a victim?	As Given
I2	N/A	The applicant was admissible	$I2Q1 \wedge I2F1$
I2Q1	f4	Is the case well founded?	As Given
I2F1	N/A	The victim suffered a disadvantage	$I2F1Q1 \wedge I2F1Q2$
I2F1Q1	f6	The case examines a fundamental aspect?	As Given
I2F1Q2	f7	Have all Domestic courts have been exhausted?	As Given
I3	N/A	Is the case fair and public	$I3F1 \wedge I3F2 \wedge I3F3 \wedge I3F4 \wedge I3F5 \wedge I3F6$
I3F1	N/A	The case was conducted in a reasonable time	$I3F1Q2 \wedge I3F1Q2$
I3F1Q1	f9	Was the case conducted in a reasonable time?	As Given
I3F1Q2	f10	Did the government cause any unreasonable delays?	As Given
I3F2	N/A	The case was independent and impartial	$I3F2Q1 \wedge I3F2Q2$
I3F2Q1	f12	The government was subjectively impartial?	As Given
I3F2Q2	f13	The government was objectively impartial?	As Given

Continue on the next page

ID	Prolog ID	Factor Text	Accepting Logic
I3F3	N/A	The case was conducted publicly and had no exceptions	$(I3F3Q1 \wedge I3F3Q2 \wedge I3F3Q3 \wedge I3F3Q4 \wedge I3F3Q5 \wedge I3F3Q6) \vee ((\neg I3F3Q1 \vee \neg I3F3Q2 \vee \neg I3F3Q3 \vee \neg I3F3Q4) \wedge (\neg I3F3Q5 \wedge \neg I3F3Q6))$
I3F3Q1	f15	If the case was public, the public wouldn't prejudice the outcome?	As Given
I3F3Q2	f16	The safety of the public wouldn't be impacted, if the case was public?	As Given
I3F3Q3	f17	Any extra privacy is not required in this case?	As Given
I3F3Q4	f18	The public would not hinder justice, if the case was public?	As Given
I3F3Q5	f19	Was the case pronounced publicly?	As Given
I3F3Q6	f20	Was the case conducted publicly?	As Given
I3F4	N/A	Equality of Arms	$I3F4Q1$
I3F4Q1	f21	Was there equality of arms?	As Given
I3F5	N/A	Access to Court	$I3F5Q1$
I3F5Q1	f41	Was the victim given appropriate access to Court?	As Given
I3F6	N/A	Legal certainty is upheld	$(I3F6Q1 \vee (\neg I3F6Q1 \wedge I3F6Q2)) \wedge \neg I3F6F1$

Continue on the next page

ID	Prolog ID	Factor Text	Accepting Logic
I3F6Q1	f42	Can the highest court be considered binding in its findings?	As Given
I3F6Q2	f43	Was the case was reopened due to new facts or a fundamental defect in fairness ?	As Given
I3F6F1	N/A	There are conflicting decisions in case law which affect the fairness of the case	$I3F6F1Q1 \wedge \neg I3F6F1Q2$
I3F6F1Q1	f44	Are there profound and long-standing differences in the case law?	As Given
I3F6F1Q2	f45	Have tools have been used to overcome any difference in case law?	As Given
I4	N/A	The applicant was presumed innocent	$I4Q1 \wedge I4Q2 \wedge I4Q3$
I4Q1	f22	Was the victim presumed innocent?	As Given
I4Q2	f38	Does the Prosecution bares the burden of proof?	As Given
I4Q3	f39	Any doubts benefited applicant?	As Given
I5	N/A	Had minimum rights	$I5F1 \wedge I5Q1 \wedge I5F2 \wedge I5F3 \wedge I5F4 \wedge I5Q2$
I5F1	N/A	The applicant was informed promptly	$I5F1Q1 \wedge I5F1Q2 \wedge I5F1Q3$
I5F1Q1	f25	Was the applicant informed in the correct language?	As Given

Continue on the next page

ID	Prolog ID	Factor Text	Accepting Logic
I5F1Q2	f26	Was the applicant given details of the case?	As Given
I5F1Q3	f27	Was the applicant told what crime they had committed?	As Given
I5Q1	f28	The applicant had time and facilities to prepare their defence?	As Given
I5F2	N/A	Opportunity to defend themselves in person	$I5F2Q3 \wedge ((I5F2Q2 \wedge \neg I5F2Q1) \vee (\neg I5F2Q1))$
I5F2Q1	f30	Has the applicant attempted to escape trial?	As Given
I5F2Q2	f31	Has the applicant waived right to defend themselves?	As Given
I5F2Q3	f40	The applicant was not prevented from accessing lawyers?	As Given
I5F3	N/A	Access to legal assistance	$I5F3Q1 \wedge I5F3Q2$
I5F3Q1	f32	Did the applicant have access to legal assistance?	As Given
I5F3Q2	f33	Did the applicant have access to free legal assistance if necessary?	As Given
I5F4	N/A	Able to examine witnesses	$I5F4Q1 \wedge I5F4Q2$
I5F4Q1	f35	Witnesses were examined under same conditions?	As Given
I5F4Q2	f37	Witnesses had a valid reason for non attendance?	As Given
I5Q2	f36	The applicant could have free access to interpreter	As Given

Table A.2: Article 6 ADF, with IDs used in the Prolog program

Appendix A.2. Admissibility ADF

ID	Factor Text	Accepting Logic
V1	Submission Recommendation	$I1 \wedge I2$
I1	Application compiles with Rule 47	$I1Q1 \wedge I1Q2 \wedge I1F1 \wedge I1F2$
I1Q1	Have you identified the state against which the application is brought to the Court (p2 of the application form)?	As Given
I1Q2	Have you ticked an appropriate box on p2 of the application form?	As Given
I1F1	All Signatures Provided	$(I1F1Q1 \wedge I1F1Q4 \wedge (I1F1Q2 \vee I1F1Q3)) \vee (\neg I1F1Q1 \wedge I1F1Q3) \vee (I1F1Q8 \wedge ((\neg I1F1Q1 \wedge I1F1Q7) \vee (I1F1Q1 \wedge I1F1Q5 \wedge (I1F1Q6 \vee I1F1Q7))))$
I1F1Q1	Do you have a legal representative?	As Given
I1F1Q2	Was the application form signed (p13) by your legal representative?	As Given
I1F1Q3	Have you (or every member of the group if an application is submitted by the group) signed the application form on page 13?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I1F1Q4	Have you (or every member of the group if an application is submitted by the group) signed the power of attorney (page 3 of the applications form)?	As Given
I1F1Q5	Has the legal representative filled out and signed the application (page 4 of the applications form)?	As Given
I1F1Q6	Has the legal representative signed the application (page 13 of the applications form)?	As Given
I1F1Q7	Has a duly authorised director or official signed the application form on page 13?	As Given
I1F1Q8	Did your duly authorised director or official filled out and signed page 4 of the application form?	As Given
I1F2	All Documentation	$I1F2Q1 \wedge I1F2Q2$
I1F2Q1	Is there a concise and legible statement of the alleged violation with the relevant arguments?	As Given
I1F2Q2	Have you attached copies of documents that relate to the decisions/measures complained of and listed them on p12?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2	Application is Admissible	$I2F1 \wedge I2F2 \wedge I2F4 \wedge I2F5 \wedge (I2F3F1 \wedge I2F3F2 \wedge \neg I2F3F3 \wedge I2F3F4 \wedge I2F3F5)$
I2F1	Valid petitioner	$I2F1Q1 \vee I2F1Q2$
I2F1Q1	Is the applicant a Physical Person or group of physical persons?	As Given
I2F1Q2	Is the applicant a legal entity?	As Given
I2F2	Victim Status	$\neg I2F2Q3 \wedge (I2F2Q1 \vee I2F2Q2 \vee I2F2Q4)$
I2F2Q1	Is the applicant a direct victim?	As Given
I2F2Q2	Is the applicant an indirect victim?	As Given
I2F2Q3	Have they lost victim status through redress?	As Given
I2F2Q4	Are you a potential victim of a violation?	As Given
I2F4	Admissible by Court's jurisdiction	$I2F4F1 \wedge I2F4F2$
I2F4F1	Ratione Personae (Government Committed)	$(I2F4F1Q1 \wedge \neg I2F4F1Q2) \vee (I2F4F1Q1 \wedge I2F4F1Q2 \wedge I2F4F1Q3)$
I2F4F1Q1	Were the alleged violations committed by a contracting party or contracting parties to the Convention?	As Given
I2F4F1Q2	Does your complaint relate to an article enshrined in one of the Protocols to the Convention?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2F4F1Q3	Has the respondent state ratified the protocol involved?	As Given
I2F4F2	Ratione loci (Committed in Government State)	$I2F4F2Q1 \vee I2F4F2Q2 \vee I2F4F2Q3 \vee I2F4F2Q4$
I2F4F2Q1	Did the violation occurred on the territory of one of the following states:Albania, Andorra, Armenia, Austria, Azerbaijan, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Montenegro, Netherlands, North Macedonia, Norway, Poland, Portugal, Republic of Moldova, Romania, Russian Federation, San Marino, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, Ukraine, United Kingdom?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2F4F2Q2	Did the violation take place on the territory occupied by one of the above-mentioned states or did the violation happen during the overseas military operation by one of the above-mentioned states?	As Given
I2F4F2Q3	Was the violation caused by diplomatic and/or consular representative?	As Given
I2F4F2Q4	Did the violation occur on board aircraft/vessels registered/flying the flag of the territory?	As Given
I2F5	Merits are Admissible	$I2F5F1 \land \neg I2F5F2F1Q1 \land I2F5F2F1Q2 \land I2F5F2F2 \land I2F5F2F3 \land I2F5F3$
I2F5F1	Scope of Court	$\neg I2F5F1Q1 \lor (I2F5F1Q1 \land I2F5F1Q2 \land I2F5F1Q3)$
I2F5F1Q1	Is the key matter of your complaint that the national court did not assess the evidence properly (fourth instance)?	As Given
I2F5F1Q2	Was the assessment of the evidence by the national Court grossly inadequate?	As Given
I2F5F1Q3	Can you prove that the assessment was grossly inadequate?	As Given
I2F5F2F1Q1	Are the facts of the case objectively impossible?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2F5F2F1Q2	Do you provide evidence of the allegations included in the application?	As Given
I2F5F2F2	No previous cases	$I2F5F2F2Q1 \vee I2F5F2F2Q2$
I2F5F2F2Q1	Has the Court ever found violations in cases similar to your case?	As Given
I2F5F2F2Q2	Does your case pose a new or unusual question before the Court?	As Given
I2F5F2F3	Domestic Remedies Exhausted Documentation	$I2F5F2F3Q1 \vee I2F5F2F3Q2$
I2F5F2F3Q1	Are there copies of documents/decisions showing that the applicant has exhausted domestic remedies and complied with 6-months rule?	As Given
I2F5F2F3Q2	Are there any valid reasons for which you could not exhaust domestic remedies?	As Given
I2F5F3	Significant Disadvantage	$I2F5F3Q1 \vee I2F5F3Q2 \vee (I2F5F3F1Q1 \wedge I2F5F3F1Q2)$
I2F5F3Q1	Have you suffered a significant financial disadvantage?	As Given
I2F5F3Q2	Was the significant disadvantage non-financial?	As Given
I2F5F3F1Q1	Does respect for human rights requires an examination of the case on the merits?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2F5F3F1Q2	Was the case duly considered by a domestic tribunal?	As Given
I2F3F1	Exhausted domestic remedies	$\neg I2F3F1F1F1 \vee (I2F3F1F1Q2 \wedge I2F3F1F1Q1) \vee (\neg I2F3F1F1Q1 \wedge I2F3F1F1Q3)$
I2F3F1F1F1	Effective Remedy Exists	$\neg I2F3F1F1F1Q1 \vee (I2F3F1F1F1Q1 \wedge I2F3F1F1F1Q2)$
I2F3F1F1F1Q1	Is there an effective remedy on the national level that the applicant has not exhausted?	As Given
I2F3F1F1F1Q2	Does the effective remedy offer reasonable chance of success?	As Given
I2F3F1F1Q1	Did the applicant do everything reasonably expected to exhaust domestic remedies?	As Given
I2F3F1F1Q2	Did the applicant complain about violation of their rights in substance?	As Given
I2F3F1F1Q3	Was there an unjustified delay for which the applicant is not responsible (for example an excessive length of criminal investigation)	As Given
I2F3F2	Complied with 6-month time limit	$I2F3F2Q1 \vee I2F3F2F1$

Continue on the next page

ID	Factor Text	Accepting Logic
I2F3F2Q1	Was the application submitted within 6 months since you were informed about the decision of the final effective remedy in your case on the domestic level?	As Given
I2F3F2F1	The violation is continuing	$\neg I2F3F2F1Q2 \wedge (\neg I2F3F2F1Q1 \vee (I2F3F2F1Q1 \wedge I2F3F2F1Q3))$
I2F3F2F1Q1	Is your violation continuing (you are arrested in an overcrowded prison you cannot get access to your property etc)?	As Given
I2F3F2F1Q2	Are there any effective remedies in your case on the national level?	As Given
I2F3F2F1Q3	Did you submit your complaint within 6 months from the moment when the violation took place?	As Given
I2F3F3	Anonymous Application	$\neg I2F3F3Q1$
I2F3F3Q1	Is the applicant identifiable?	As Given
I2F3F4	Substantially Different	$I2F3F4F1 \wedge \neg I2F3F4F2$
I2F3F4F1	The same application has not been submitted to the ECHR	$I2F3F4F1Q3 \vee \neg I2F3F4F1Q2$
I2F3F4F1Q1	Have you submitted the case to the European Court of Human Rights previously?	As Given
I2F3F4F1Q2	Was the complaint in the previous case the same?	As Given

Continue on the next page

ID	Factor Text	Accepting Logic
I2F3F4F1Q3	Are there significantly new facts or developments in your case?	As Given
I2F3F4F2	The same application been submitted to another international body	$I2F3F4F2Q1 \wedge I2F3F4F2Q2 \wedge I2F3F4F2Q3$
I2F3F4F2Q1	Is there a case submitted to the international body that relies on the same facts?	As Given
I2F3F4F2Q2	Was the case to another body submitted by the same applicant?	As Given
I2F3F4F2Q3	Was the complain the same?	As Given
I2F3F5	Not abused the right of application	$\neg I2F3F5F1 \wedge \neg I2F3F5F2 \wedge \neg I2F3F5F3$
I2F3F5F1	Misleading	$I2F3F5F1Q1 \vee I2F3F5F1Q2 \vee I2F3F5F1Q3 \vee I2F3F5F1Q4$
I2F3F5F1Q1	Was the application submitted under a false identity?	As Given
I2F3F5F1Q2	Were any documents falsified?	As Given
I2F3F5F1Q3	Has the court been misled by the applicant regarding the core of the case?	As Given
I2F3F5F1Q4	Has any new important information not been disclosed by the applicant?	As Given
I2F3F5F2	Offensive language	$I2F3F5F2Q1 \wedge I2F3F5F2Q2$

Continue on the next page

ID	Factor Text	Accepting Logic
I2F3F5F2Q1	Have you used any offensive language used in the application form or in any accompanied documents?	As Given
I2F3F5F2Q2	Was the offensive language outside the bounds of normal civil and legitimate criticism?	As Given
I2F3F5F3	Devoid of real purpose	$I2F3F5F3Q1$
I2F3F5F3Q1	Have you lodged similar applications that were declared by the Court manifestly ill-founded?	As Given

Table A.3: Admissibility ADF

Appendix A.3. No Violation Example

```
The applicant is the victim
        The case is well founded
        The case does examine a fundamental part of human
            rights act
        All 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 would not prejudice outcome of case
        Safety of the public would not be impacted by the case
             being publicly pronounced
```

Privacy is not required to deliver justice
The public would not hinder delivery of justice
The case was pronounced publicly
The case was conducted publicly
The case was required to be conducted publicly, and was
There was a fair balance between parties, and had equality of arms
Given appropriate access to a court
The highest court was considered binding in its findings
The case was not reopened due to new facts or a fundamental defect in fairness
There are not profound and long–standing differences in the case law
Tools have not been used to overcome any difference in case law
There are not conflicting decisions in case law which affect the fairness of the case
Principle of Legal Certainty is upheld
The case was 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 mount 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 no violation of Article 6
```

Appendix A.4. Test Cases

Firstly we present in Table A.4 the evaluation results for the 10 cases that were also used in previous work [69]. For clarity, the breakdown of these cases includes 5 violation cases, 4 non violation cases, and an inadmissible case.

Case Name	Expected Prediction	Actual Prediction
MARGUŠ v. CROATIA	No violation	No violation
CARDOT v. FRANCE	Inadmissable	Inadmissable
ABDULLAYEV v. RUSSIA	Violation	Violation
ZARKOV v. SERBIA	Violation	Violation
MOSER v. AUSTRIA	Violation	Violation
CHAPMAN v. THE UNITED KINGDOM	No violation	No violation
KHANUSTARANOV v. RUSSIA	Violation	Violation
STOILKOVSKA v. THE FORMER YUGOSLAV REPUBLIC OF MACEDONIA	Violation	Violation
UŽKURĖLIENĖ AND OTHERS v. LITHUANIA	No violation	No violation
T.P. AND K.M. v. THE UNITED KINGDOM	No violation	No violation

Table A.4: Results for the same cases that were used in [69]

We can see that our extended Article 6 Prolog implementation reported in this paper is able to achieve the same 100% accuracy as achieved in our earlier work [69] that considered these 10 cases.

We now give results for a new evaluation conducted on 30 additional cases, all of which are included in the Aletra's *et al.* corpus [4]. In Table A.5 we show the results for the 15 cases which were found to have a violation and in Table A.6 we show the results on the 15 no violation cases. One violation case was incorrect: this is discussed in Section 6.2, and results from a misdescription of the case in terms of factors, rather than a flaw in the model.

Case Name	Expected Prediction	Actual Prediction
BEZBORODOV v. RUSSIA	Violation	Violation
FETAOVSKI v. THE FORMER YUGOSLAV REPUBLIC OF MACEDONIA	Violation	Violation
GRIDAN AND OTHERS v. ROMANIA	Violation	Violation
HEYDAROVA v. AZERBAIJAN	Violation	Violation
JOVANOVIĆ AND OTHERS v. SERBIA	Violation	Violation
KARANOVIĆ v. BOSNIA AND HERZEGOVINA	Violation	No Violation
KHANUSTARANOV v. RUSSIA	Violation	Violation
KOKSHAROVA v. RUSSIA	Violation	Violation
TRANČÍKOVÁ v. SLOVAKIA	Violation	Violation
MOMIĆ AND OTHERS v. BOSNIA AND HERZEGOVINA	Violation	Violation
ŠEKEROVIĆ AND PAŠALIĆ v. BOSNIA AND HERZEGOVINA	Violation	Violation
PREMOVIĆ v. SERBIA	Violation	Violation
VELINOV v. THE FORMER YUGOSLAV REPUBLIC OF MACEDONIA	Violation	Violation
VELSKAYA v. RUSSIA	Violation	Violation
STOJILKOVIĆ AND OTHERS v. SERBIA	Violation	Violation

Table A.5: Prediction results on the 15 violation cases, with 14 out of 15 cases returning a correct prediction.

Case Name	Expected Prediction	Actual Prediction
ALBU AND OTHERS v. ROMANIA	No violation	No violation
ALCHAGIN v. RUSSIA	No violation	No violation
ALKES v. TURKEY (No. 2)	No violation	No violation
BALOGH v. HUNGARY	No violation	No violation
BYKOV v. RUSSIA	No violation	No violation
DALLOS v. HUNGARY	No violation	No violation
DEMEBUKOV v. BULGARIA	No violation	No violation
DONOHOE v. IRELAND	No violation	No violation
ELEZI v. GERMANY	No violation	No violation
FADIN v. RUSSIA	No violation	No violation
HASAN AND CHAUSH v. BULGARIA	No violation	No violation
JANE SMITH v. THE UNITED KINGDOM	No violation	No violation
JUSSILA v. FINLAND	No violation	No violation
KIENAST v. AUSTRIA	No violation	No violation
MARCELLO VIOLA v. ITALY	No violation	No violation

Table A.6: Prediction results on the 15 no violation cases, with all returning a correct prediction.

Appendix B. Questionnaires

Appendix B.1. Pilot Usability Study Questionnaire

1. (Functionality) **Does the program have a reasonable response time?**
    - The program responds with no significant delay
    - The program is occasionally slow to respond

- The program is frequently slow to respond
- The program is too slow to be usable

2. (Functionality) **Did the program run to completion without any interruptions?**
    - Ran to completion every time
    - Ran to completion most of the time
    - Ran to completion some of the time
    - The program never ran to completion
3. (Usability) **How easy was the program to use?**
    - Extremely easy, no issues encountered in using the program
    - The program is as usable as most computer programs.
    - The program is harder to use than most computer programs.
    - The program is very hard to use due to constant issues.
4. (Usability) **How intuitive was the program to start using?**
    - The program is extremely easy to start using; how to interact with it was immediately obvious
    - Using the program is obvious after a small amount of training
    - Using the program is not immediately obvious after a small amount of training
    - The program is hard to start using; there would need to be extensive training to be able to use it.
5. (Explainability) **How effective was the explanation given for describing the program's decisions?**
    - The program's explanations were clear and appropriate
    - The program justified the decisions made well enough
    - The program's explanations were not fully clear
    - The program's explanations were unclear and confusing
6. (Explainability) **How easy was the information to parse?**

- The explanations provided were very easy to parse
- The explanations provided were easy to parse
- The explanations provided were hard to parse
- The explanations provided were very hard to parse

7. (Usefulness) **How useful would you find this program for assisting you in your work?**
   - The program is extremely useful and would use as is
   - The program is useful, though some changes are needed
   - The program is not useful, many changes are needed
   - The program is not useful at all, changes would not change the usefulness
8. (Usefulness) **Generally how useful would additional technology be for assisting with legal work?**
   - Technology is needed in the law domain to rapidly to improve over current service levels.
   - Technology has a place in the legal domain, but needs careful development
   - Technology has limited use to the legal domain
   - Technology is not useful to the legal domain and should not be incorporated.
9. (Questions about Questions) **How clear were the questions that you answered within the program?**
   - I understood all the questions
   - I understood most of the questions
   - I didn't understand many of the questions.
   - I understood none of the questions:

   (a) (Questions about Questions) **Which questions were not understood?**
10. (Questions about Questions) **How much time would you save if you used a fully functional program for your work on deciding on the admissibility of cases?**
    - A significant amount of time would be saved

- Some time would be saved
- No time would be saved
- Extra time would be needed to use the program

11. (Questions about Questions) **Does the program reflect how you decide on the admissibility of cases that you process?**
    - The program exactly reflects how I would process admissibility
    - The program mostly reflects how I would process admissibility
    - The program does not reflect the current process but does capture some aspects of decision making on admissibility
    - The program does not at all reflect how I would process admissibility cases

Appendix B.2. Wider Survey Questionnaire

1. ***Please click Yes to confirm that you have given consent to process your responses in this questionnaire.**
    - Yes
    - No
2. **Does the program have a reasonable response time?**
    - The program responds with no significant delay
    - The program is occasionally slow to respond
    - The program is frequently slow to respond
    - The program is too slow to be usable
3. **Did the program run to completion without any interruptions?**
    - Ran to completion every time
    - Ran to completion most of the time
    - Ran to completion some of the time
    - The program never ran to completion
4. **How easy was the program to use?**
    - Extremely easy, no issues encountered in using the program

- The program is as usable as most computer programs.
- The program is harder to use than most computer programs.
- The program is very hard to use due to constant issues.

5. **How intuitive was the program to start using?**
    - The program is extremely easy to start using; how to interact with it was immediately obvious
    - Using the program is obvious after a small amount of training
    - Using the program is not immediately obvious after a small amount of training
    - The program is hard to start using; there would need to be extensive training to be able to use it.
6. **How effective was the explanation given for describing the program's decisions?**
    - The program's explanations were clear and appropriate
    - The program justified the decisions made well enough
    - The program's explanations were not fully clear
    - The program's explanations were unclear and confusing
7. **How easy was the information to parse?**
    - The explanations provided were very easy to parse
    - The explanations provided were easy to parse
    - The explanations provided were hard to parse
    - The explanations provided were very hard to parse
8. **How useful would you find this program for assisting you in your work?**
    - The program is extremely useful and would use as is
    - The program is useful, though some changes are needed
    - The program is not useful, many changes are needed
    - The program is not useful at all, changes would not change the usefulness

9. **Generally how useful would additional technology be for assisting with legal work?**
    - Technology is needed in the law domain to rapidly to improve over current service levels.
    - Technology has a place in the legal domain, but needs careful development
    - Technology has limited use to the legal domain
    - Technology is not useful to the legal domain and should not be incorporated.
10. **How clear were the questions that you answered within the program?**
    - I understood all the questions
    - I understood most of the questions
    - I didn't understand many of the questions.
    - I understood none of the questions:
11. **Which questions were not understood?**
12. **How much time would you save if you used a fully functional program for your work on deciding on the admissibility of cases?**
    - A significant amount of time would be saved
    - Some time would be saved
    - No time would be saved
    - Extra time would be needed to use the program
13. **Does the program reflect how you decide on the admissibility of cases that you process?**
    - The program exactly reflects how I would process admissibility
    - The program mostly reflects how I would process admissibility
    - The program does not reflect the current process but does capture some aspects of decision making on admissibility
    - The program does not at all reflect how I would process admissibility cases
14. ***Do you have any other comments/thoughts you would like to provide, either to expand on your above responses, or in addition to these?**