Explaining Factor Ascription

Jack Mumford, Katie Atkinson and Trevor Bench-Capon

Department of Computer Science,
University of Liverpool, UK

Abstract. Explanation and justification of legal decisions has become a highly relevant topic in light of the explosion of interest in the use of machine learning (ML) approaches to predict legal decisions. Current suggestions are to use the established factor based explanations developed in AI and Law as the basis for explaining such programs. We, however, identify factor ascription as an important aspect of explanation of case outcomes not currently covered, and argue that explanations must also include this aspect. Finally, we outline our proposal for a hybrid system approach that combines ML and Abstract Dialectical Framework (ADF) layers to engender an explainable process.

Keywords. reasoning with cases, explanation, factor ascription.

1. Introduction

In recent years there has been an explosion of interest in the use of machine learning (ML) to predict legal decisions (e.g. [14]). A major weaknesses of these approaches is, however, that they are unable to explain their reasoning in an acceptable manner. Traditional explanations of ML such as listing or highlighting the most influential words in the texts have been shown to be unhelpful [9] because they are difficult to relate to the relevant law. Moreover there are good reasons why any such explanation would be inappropriate in a legal context [8]. The right to explanation means that the explanation must be capable of persuading the losing party, and providing a justification which can withstand an appeal. It need not be an explanation of how the decision was in fact reached, but must explain why the decision represents the proper application of the law[1].

In order to explain the predicted outcomes in appropriate terms, researchers have turned to the extensive body of work on explanation developed in AI and Law [5]. In particular the type of explanation advocated has been based on the use of factors, as developed in CATO [2], e.g [9], [16] and [6]. Specifically, using the set of argumentation schemes designed to capture the reasoning of CATO from [17] is advocated in [16].

---

[1] A popular caricature of legal realism (e.g. [12]) says that the law is what the judge had for breakfast. This may indeed explain *how* the judge reached the decision, but the opinion must contain an explanation of *why* this decision is a justifiable application of the law.

Factors as introduced in [2] are stereotypical patterns of facts which are legally significant in that they provide a reason to find for one side or the other. They represent a generalisation from the facts of particular cases so that they can be applicable to a number of cases: for example *plaintiff pursuing livelihood* generalises the facts of several property law cases involving hunting, shooting and fishing. In [2], the factors are organised into a factor hierarchy with issues at the upper levels, abstract factors in the middle layers and the (base level) factors as the leaves. CATO organises its explanation into a series of issues. The resolution of issues can be explained in terms of the factors which provide reasons for the winning party or the burden of proof for that issue. Where there are factors for both the plaintiff and the defendant, the reasons preferred are justified in terms of a precedent case exhibiting this preference. Sometimes factors may be used to cancel or substitute for other factors as described in [17]. Once the issues have been resolved, the decision follows logically according to a logical model of issues [3], which also serves to associate the factors with particular issues.

Explanation now begins with a summary in terms of the issues. In CATO's domain, US Trade Secrets Law, the plaintiff must establish that the information is a trade secret and that it was misappropriated. Misappropriation requires either the use of improper means or breach of confidence. Thus in *Mason v. Jack Daniels* the explanation begins with: *Plaintiff should win. Plaintiff's information is a trade secret, a confidential relationship existed between plaintiff and defendant, and defendant acquired plaintiff's information through improper means.*

Each issue is then explained in terms of the factors and the relevant precedents. For example, the factors relating to confidential relationship were F1 (*DisclosureInNegotiations*) and F21 (*KnewInfoConfidential*). That the latter is a stronger reason than the former was established in *Forest Labs, Inc. v. Formulations, Inc*. So the explanation of the issue is: *A confidential relationship exists because although the information was disclosed in negotiations, the defendant knew that the information was confidential (Forest Labs, Inc. v. Formulations, Inc).*

Explanations using argument schemes in [16] are different: they produce a three layer tree of argument schemes. The top layer cites the most on point precedent, the second layer attacks this argument with distinctions and counter examples, while the third rebuts the counter examples through a series of transformations and rebuts the distinctions by cancellation and substitution. No use is made of issues in [16], but [6] suggests that the schemes should be applied at the issue level rather that at the case level. Explanation using factors provides a good explanation of cases like *Mason*, but for other cases they are less satisfactory.

2. Explaining Factor Ascription

Consider the Trade Secrets case of *Arco Industries Corp. v. Chemcast Corp*. In that case we have three factors: F10 (*InfoDisclosedToOutsiders*), F16 (*InfoReverseEngineerable*) and F20 (*InfoKnownToCompetitors*): all favour the defendant and establish that the information was not a trade secret. The explanation is thus: *the defendant should win because the information is not a trade secret. It had been disclosed to outsiders, was known to competitors and was reverse engineerable.*

Will this satisfy the plaintiff? Arco had in fact argued that the information *was* a trade secret because it was covered by a patent. They are therefore essentially arguing that two pro-plaintiff factors, F15 (UniqueProduct) and F18 (IdenticalProduct) are present. These factors might have been sufficient to establish that the information was indeed a trade secret and that the information had been used. If so, the plaintiff may well have established F21 (*KnewInfoConfidential*) and thus won the case. So the explanation needs to cover the reasons why the arguments based on these factors were rejected. The decision in fact includes a detailed discussion of the patent specification and the product (a grommet) produced by the defendant. It concludes: *The specifications describe the "recess" as an indentation below the planar surface of the grommet which in turn lies below the peripheral sealing ridge. The accused grommet does not have such a recess.*

This is what excludes F15 and F18 and makes discussion of confidentiality unnecessary. It is the general notion of a grommet that is known in the industry and it appears that the unique feature of Arco's grommet was not used by Chemcast. The explanation needed by the plaintiff is not how the factors ascribed satisfy the required issues, but why the claim that other factors were present was rejected. Crucially, the explanation in terms of factors fails to answer the plaintiff's question: *why was the information not protected by the patent?*

In Arco it was the absence of factors that needed to be explained, but sometimes the presence of a factor needs to be explained. Consider *A. H. Emery Co. v. Marcan Products Corporation*. In that case the information had been learned by the defendants while they were employed by the plaintiff. The defendants had not signed any non disclosure agreements, and so they denied that they had breached a confidential relationship. Nonetheless F21 (KnewInfoConfidential) was held to apply, and so a confidential relationship was held to exist. The point was that the information had been acquired while the defendants were employed by the plaintiff and they knew the information to be confidential and at the time the information was acquired they owed a duty of fidelity to their employer.

3. Discussion of Challenges

Reasoning with legal cases is a two stage process [7]: first factors are ascribed and then the balance of factors is determined to reach a decision. Both aspects require explanation. *Arco* and *Emery* show the need to be able to offer explanations not only of the balance of competing factors, but also the presence and absence of particular factors. This could be delivered by extending the dialogue of [6] to ask *WHY?* of any factor used to explain an issue and *WHY NOT Fn?* of any factor not mentioned in the explanation of an issue.

Despite the attention in AI and Law paid to explaining precedential reasoning, there has been little or no work on explaining why the factors are present or absent. This is because most research since HYPO [18] has taken the factors as given. HYPO's dimensions give a clue as to how we might explain certain factor ascriptions. The ranges on these dimensions in which factors are applicable are defined in precedents [7]. So we can explain the ascription of such a factor in terms of these precedents. For example if we have a precedent (*PrecL*) establishing that

an absence of 15 months is a *long stay*, we can explain the ascription of this factor to a new case (CaseN) with an absence of 17 months by saying *long stay applies to CaseN because the absence was greater than 15 months (PrecL)*.

This kind of explanation is promising for factors which can be seen as ranges on well ordered dimensions, but does not seem applicable to the kind of detailed consideration of very particular facts that we saw in *Arco* and *Emery*. Such cases may involve analogy [4], or some kind of common sense ontology. This suggests that the key role for ML is not the prediction of outcomes, but the identification of the factors as in [3], [19] and [9].

How does the issue of explaining ascription relate to ML approaches? If we follow [9] we must accept that, if the explanation is to be given in acceptably legal terms, the ML system will need to learn to ascribe factors as well as predict outcomes. Although the standard ML explanations of outcomes are unsatisfactory, there is a considerable gap between facts and outcomes, requiring reasoning through factors and issues. There is no such conceptual gap between facts and factors, and so it may be that the explanation of the ascription of factors using ML is more satisfactory. This is something that requires empirical investigation.

4. Next Steps

Our approach to producing explainable case predictions is to separate the process within a hybrid system in accordance with the two stages outlined in the previous section. The first stage, factor ascription, will be addressed via ML natural language processing (NLP). The second stage, reaching a decision, will be addressed via the balancing of factors within a pre-determined non-cyclic Abstract Dialectical Framework (ADF) [10] that has been derived with expert knowledge to capture the factor-based reasoning of a legal domain, as demonstrated in [1]. If we accept the argument of [15] then domain expertise is of paramount importance when establishing an appropriate ADF, rendering data-driven approaches less effective at the level of factors and above. As such, our hybrid system is initially poised to only adjust the architecture of the NLP layer; the ADF layer will not be changed from its initial state as rendered by expert judgement.

For the first stage of the process, our intention is to use a state-of-the-art Hierarchical BERT model, similar to the approach taken in [13] only not used for determining the case outcome, but rather for factor ascription. We propose to use a Hierarchical BERT model due to the proffered combination of impressive classification performance and the sentence-level attention weights that could sufficiently express the relevant facts that explain a given factor's ascription or non-ascription. The model takes a natural language description of a given case as input and outputs a binary classification of 'ascribed' or 'not ascribed' for each base-level factor in the ADF. The second stage of the process will use the expert-derived ADF to produce a decision via the reasoning steps following from the base-level factor ascription input from the NLP layer.

The data-driven learning phase will not amend the ADF, as previously stated, but still passes errors through the levels in the framework down to the base-level factors for use with the NLP classification task for factor ascription. The full de-

scription of the algorithms we define for error propagation through the ADF will be set out in future work. However, from a high level perspective, the algorithms will function by creating a graphical scaffold of the ADF in which each node is a linearly separable function where children nodes are only capable of attack (that is, each child node implies the contradiction of the parent node), in order to facilitate computationally tractable error propagation. Our initial attention will be on legal domains represented by ADFs with Boolean acceptance conditions. Given the discontinuous Heaviside step function that governs the Boolean acceptance of any given node, backpropagation is not appropriate for error propagation in general. Instead, when a case is input to the hybrid system that results in the wrong decision, errors will be propagated backwards through the levels of the graphical scaffold that give each factor a tuple of weights (*ascribed, not ascribed*), where the value of *ascribed* (alternatively *not ascribed*) is scaled by the product of its parent's value for *not ascribed* (alternatively *ascribed*) and the proportion of combinations in which the node is ascribed (alternatively not ascribed) that would cause the parent node to not be ascribed. This iterative process down through the levels of the scaffolding will begin with the root node representing the decision, which will have a tuple of: (0, 0) if there is no error, (1, 0) if it should be ascribed, or (0, 1) otherwise. The scaffolding will work in accordance with abstract argumentation stable semantics [11], since no cycles are permitted and the ascription of any parent node is determined by the ascription of its children. The algorithms are thus intended to be used to render a tuple of weights for the base-level factors which can then be used to determine the proportion of the classification tasks assigned to the NLP layer. For example, if a factor F1 has a tuple (0.1, 0.5), then in the next training epoch we would run five times as many classification tasks for F1 being not ascribed as we would for F1 being ascribed. This example also illustrates the benefit of partitioning the ADF layer to accommodate only Boolean nodes, since most NLP tasks involve classification and not regression. Future work will look at extensions to encompass ADFs that include non-Boolean nodes, but further consideration would need to be given as to how to extract continuous valued data points from text.

5. Concluding Remarks

In summary, in this short paper we have identified the need for the explanations of legal decisions to go beyond factors and preferences between them, and explain the ascription and non-ascription of the factors themselves. This sort of explanation is as yet largely uninvestigated in AI and Law, which has taken the factors in a case as given. We are currently engaged in research to establish the nature of such explanations through a hybrid approach, with the first stage in the process being addressed through ML and the second stage through reasoning via the medium of an ADF. Our immediate focus is on formal articulation and application of the aforementioned algorithms and Hierarchical BERT model to enable these tasks.

References

[1] L Al-Abdulkarim, K Atkinson, and T Bench-Capon. Statement types in legal argument. In *Proceedings of JURIX 2016*, pages 3–12. IOS Press, 2016.

[2] V. Aleven. *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh, 1997.

[3] K Ashley and S Brüninghaus. Automatically classifying case texts and predicting outcomes. *AI and Law*, 17(2):125–165, 2009.

[4] K Atkinson and T Bench-Capon. Reasoning with legal cases: Analogy or rule application? In *Proceedings of the 17th ICAIL*, pages 12–21. ACM, 2019.

[5] K Atkinson, T Bench-Capon, and D Bollegala. Explanation in AI and Law: Past, present and future. *Artificial Intelligence*, 289:103387, 2020.

[6] T Bench-Capon. Using issues to explain legal decisions. *XAILA worhshop at ICAIL 2021. arXiv preprint arXiv:2106.14688*, 2021.

[7] T Bench-Capon and K Atkinson. Precedential constraint: The role of issues. In *Proceedings of the 18th ICAIL*, pages 12–21. ACM, 2021.

[8] F Bex and H Prakken. On the relevance of algorithmic decision predictors for judicial decision making. In *Proceedings of the 18th ICAIL*, pages 175–179. ACM, 2021.

[9] L K Branting, C Pfeifer, B Brown, L Ferro, J Aberdeen, B Weiss, M Pfaff, and B Liao. Scalable and explainable legal prediction. *AI and Law*, 29(2):213–238, 2021.

[10] G Brewka, S Ellmauthaler, H Strass, J Wallner, and P Woltran. Abstract dialectical frameworks revisited. In *Proceedings of the Twenty-Third IJCAI*, pages 803–809. AAAI Press, 2013.

[11] 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):321–357, 1995.

[12] A Kozinski. What I ate for breakfast and other mysteries of judicial decision making. *Loyola of Los Angeles Law Review*, 26:993, 1992.

[13] M Medvedeva, A Üstun, X Xu, M Vols, and M Wieling. Automatic judgement forecasting for pending applications of the european court of human rights. In *Proceedings of ASAIL 2021*, 2021.

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

[15] J Mumford, K Atkinson, and T Bench-Capon. Machine learning and legal argument. In *Proceedings of the 21st CNMA Workshop, CEUR Workshop Proceedings*, volume 2937, pages 47–56, 2021.

[16] H Prakken and Ratsma R. A top-level model of case-based argumentation for explanation: formalisation and experiments. *Argument and Computation*, Available On-Line. 2021.

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

[18] E Rissland and K Ashley. A case-based system for Trade Secrets law. In *Proceedings of the 1st ICAIL*, pages 60–66. ACM, 1987.

[19] A Wyner and W Peters. Lexical semantics and expert legal knowledge: towards the identification of legal case factors. In *Proceedings of Jurix*, pages 127–136. IOS, 2010.