The Ideal Audience and Artificial Intelligence and Law

Trevor Bench-Capon

Department of Computer Science
The University of Liverpool, Liverpool, UK.
E-mail: tbc@csc.liv.ac.uk

Abstract

In this talk I will reflect on the notion of an ideal audience and its implications for AI and Law. After introducing the notion, I consider how it can be used as an analytical tool for thinking about AI and Law systems and the practice of AI and Law.

1. Introduction

In this talk I will discuss an important, though somewhat overlooked, aspect of argument: what is it that makes one argument more convincing than another? To get some purchase on this question I shall use the notion of the *ideal audience,* drawing on work by Perelman, and some more recent work by George Christie. I shall then show how this notion can be used as an analytical tool for reflecting on AI and Law systems the practice of AI and Law.

Much excellent work has been done to suggest ways in which arguments can be generated, analysed and used in legal reasoning. What it has not so far been addressed with much success is the question of what, given two conflicting arguments, makes one argument more convincing than another. In Bench-Capon and Sergot (1988), the concluding sentence was:

> "In the longer term, we hope to pursue what we have identified as a critical requirement: a representation in computer-intelligible terms of what it is that makes a legal argument persuasive."

2. Perelman and the Universal Audience

Perhaps the central idea in Perelman's work is that an argument, indeed all speech, is directed towards an audience. But the audience towards which argument is directed is essentially a *construct* of the speaker: If we take this seriously we can see that whether an argument is accepted is not something that can be determined by considering the intrinsic qualities of the argument: rather it is a function both of the argument, and the audience to which it is addressed. But this starts to look dangerously *ad hominem*: while I may convince a particular group of people by correctly understanding and exploiting their prejudices and misconceptions, such a victory does not guarantee that my claim has any objective right to acceptance. Perelman, of course, was well aware of this problem:

> "Argumentation aimed exclusively at a particular audience has the drawback that the speaker, by the very fact of adapting to the views of his listeners, might rely on arguments that are foreign or even directly opposed to what is acceptable by persons other than those he is presently addressing." (Perelman and Olbrechts-Tyteca 1969, p31).

For this reason opinions that enjoy unanimous approval are considered stronger than those accepted by particular audiences only. But this casts no light on whether an argument *should* be accepted, only whether it is, as a matter of fact, widely accepted. Many arguments aspire to more than this: to reflect this Perelman introduces the notion of the *universal* audience. This universal audience is not simply an aggregation of all particular audiences: acceptance by the universal audience refers

> "not to an experimentally proven fact, but to a universality *imagined* by the speaker, to the agreement of an audience which *should* be universal, since for *legitimate* reasons, we need not take into consideration those which are not part of it." (Perelman and Olbrechts-Tyteca 1969, p31, *italics mine*).

Those who address the universal audience

> "think that all who understand their reasons will have to accept their conclusions. *The agreement of a universal audience is thus a matter, not of fact, but of right*". (Perelman and Olbrechts-Tyteca 1969, p31, *italics theirs*).

At this point we might be tempted to believe that the universal audience will accept only claims that can be logically proven. If we are doing mathematics this may even be correct. But this is not Perelman's intention: that would limit the realm of what can be argued before the universal audience too drastically.

> "Everyone constitutes the universal audience from what he knows of his fellow men, in such a way as to transcend the few oppositions he is aware of. Each individual, each culture, has thus its own conception of the universal audience. The study of these variations would be very instructive, as we would learn from it what men, at different times in history, have regarded as *real*, *true*, and *objectively valid*". (Perelman and Olbrechts-Tyteca 1969, p31, *italics theirs*).

Perelman does not perform such a study. We will, however, start an exploration of these variations in the next section.

3. Christie and The Ideal Audience

Christie (2000) shifts from Perelman's notion of a *universal* audience to one of an *ideal* audience. This emphasises the constructive nature of the concept: it is not really universal, but rather seen as universal from within a particular culture.

> "Some aspects of the ideal audience that manifest themselves in legal argumentation seem to have truly universal acceptance; others, although thought of as universal, seem more closely tied to particular cultures." (Christie 2000. p 41).

What are these differences? Christie begins with examples relating to the interpretation of statutes and treaties. He contrasts the narrow, cramped and particularistic view of interpretation of some courts (US in his examples), with the more expansive interpretation of others (Israeli and German).

Christie attempts to account for these differences in terms of competing visions of the way the state should be organised, and competing visions of the purpose of the state. He contrasts states which exercise authority through a hierarchically organised set of officials, with those which exercise authority through a large number of co-ordinate officials. Also he contrasts the "reactive" state which exists to resolve conflicts, and the "activist" state, which has its own notion of the good which it attempts to promote. Moreover, even accepting that the law exists to promote the common good, should we see *the common good* as composed of individual goods, or as an independent entity?

Christie then identifies another tension:

> "that between the intellectual urge to greater generality, and the cautionary counsel of experience to focus as much as possible on the particular. The ideal or universal audiences to which legal argument is directed are constantly torn by these conflicting impulses". (p108).

A third source of difference between ideal audiences is the amount of freedom that they conceive a decision maker such as a judge should have. Here again we have a tension: there are good reasons both to seek to broaden discretion, and to narrow it. At any given time a given legal system will prefer a particular position between them.

Finally we need to consider consistency - the degree to which diverse outcomes are tolerated. Christie argues that a civil-law tradition is less likely to accept that there are many reasonable answers to a problem, and to be less tolerant of inconsistency. This is important:

> "the accepted style of legal reasoning has a profound effect on the substance of the decisions the courts will make. Style is not a matter of mere form. It sometimes has a decisive influence on the decision of cases." (p 192).

There are, then, several specifically legal factors which characterise the ideal audience for legal argument. These are a series of tensions: between narrow and broad construction; between a reactive and an activist state; between the common good as an aggregation of individual goods and as an independent notion; between the particular and the general; between wide discretion and limited discretion; and between consistency and diversity. The position between these extremes ascribed to the ideal audience will vary across cultures and within cultures over time. The lawyer constructing an argument gauges the position of the ideal audience to be addressed by being part of the relevant culture.

4. The Ideal Audience and Values

There is, however, another aspect to the ideal audience which emerges from Christie's examples - that the ideal audience embodies a set of values, and of preferences between those values when they conflict.

In *Gregg v. Georgia* (428 US 153 1976), in which the Supreme Court upheld a sentence of death, there were two dissenting arguments, from Justices Brennan and Marshall. Marshall's dissent is quoted:

> "In *Furman* [an earlier case] I observed that the American people are largely unaware of the information critical to a judgement on the morality of the death penalty, and concluded

> that if they were better informed they would consider it shocking, unjust and unacceptable … the death penalty is unconstitutional because it is excessive. An excessive penalty is invalid under the Cruel and Unusual Punishments Clause even though popular sentiment may favour it" (quoted on p 22-3 of Christie 2000).

This dissent is clearly addressed to an ideal audience, one which is "better informed" than the audience that will receive it. It is addressed in a sense to a future audience, which will hold more enlightened values. It accepts that it will fail to convince its immediate auditors, but seeks to promote a change in their values, so that it will be accepted at a later date.

Another example concerns pre-Civil war cases in the US in which judges were called upon to enforce fugitive slave laws. Some refused to enforce them, offering legally suspect reasons, but appealing to the rejection of slavery they ascribed to the ideal audience. Others did enforce them, not because they condoned slavery, but because they "felt that preserving the integrity of the federal union, which was threatened with secession by the slave states, *was the greater good*" (p 45, *italics mine*).

Thus in addition to embodying the various positions that can be taken as to the organisation and purpose of the state, and various preferences for styles of legal reasoning, the ideal audience must be seen also as embodying a set of values and a system of preferences amongst those values.

5. The Ideal Audience and AI and Law Systems

What are the implications of a study of the ideal audience for AI and Law? How is the ideal audience represented in a computer system? There are several perspectives we can take, but first we should notice that cognisance of the ideal audience is not usually explicit in AI and Law systems. Rather the system builder's construct of the ideal audience tends to be built into the system, incorporated in the very principles according to which the system is constructed. The system does not construct the ideal audience, rather the system is constructed according to the system builder's presupposition of an ideal audience. Nevertheless, it is instructive to use the notion of the ideal audience to analyse some issues in AI and Law.

First we can use the notion of ideal audience to account for differences between systems and approaches in AI and Law. Let us remember that the nature of the ideal audience is *presupposed* within particular legal cultures. Thus people building AI and Law systems will have a tendency to picture the ideal audience in terms of the culture of which they are a part.

There used to be a lively debate in AI and Law between proponents of a case based approach and proponents of a rule based approach. Panels addressed this topic at both the First and Third International Conferences on AI and Law, held in 1987 and 1991. Without too much of a caricature, we can relate these approaches to some of the positions identified as characteristics of the ideal audience above. Preferring the particular over the general gives a tendency towards a case based approach. Case based approaches can be seen as reactive and rule based approaches as activist. Rule based systems tend to promote consistency, and minimise discretion. In sum, case base approaches might be expected to emerge from a common law culture, and rule based systems from a civil-law culture. This is borne out by the fact that case based approaches were more popular in the US, and rule based approaches were more popular in Europe. To meet objections that rule based approaches were also popular in the UK, one may point to the fact that much of the UK work concerned administrative welfare law, where the activist state, lay adjudication, and the need for consistency, push it towards the civil-law position. In the late 1980s and early 1990s, when AI and Law was only beginning to develop an international community, people would be prepared to argue that one position or the other was *right*. Such argument was based on the feeling that the ideal audience was the one which adopted the position of the ideal audience from the culture from which its proponent came. The debate has now died down, and there is far more of a tendency to see the approaches as complementary, or suitable for use in different circumstances, than as either right or wrong absolutely. In part this is a matter of recognising that there is no single, universal ideal audience, and in part because the AI and Law community has developed its own construct of the ideal audience, which is capable of recognising merit in both approaches. There have been gains and losses in this development. Gains stem from the ability to learn from other approaches (Prakken and Sartor (1998), in which they model a classic case based system, HYPO, in their preferred rule based terms, is an excellent example), and a more receptive climate for discussion within the AI and Law community. Losses have come when we attempt to return results from AI and Law to the legal community. Because the systems are no longer so tied to the ideal audiences presupposed by particular communities, they speak less clearly to them, and assumptions embodied in the system which are not those of the legal community mean that they can be seen as actually *wrong* by those communities. Thus the relaxation of the links between AI systems and particular constructs of the ideal audience have facilitated progress in AI and Law, but militated against acceptance of them by legal communities.

From this we should expect the systems which gain most acceptance outside of AI and Law to be those which both have a clear conception of the ideal audience of the legal community at which they are targeted, and which most closely reflect this audience.

We can illustrate this by reference to several systems. First consider the fielded systems in the Netherlands, which tend to be rather straightforward rule based expert systems dealing with administrative law, especially welfare benefits, and based on shells such as TESSEC (de Bakker and Wassink 1991). These were deployed in a country which has a civil-law tradition, and in an application domain where emphasis is on consistency and accuracy. Moreover the applications were developed in close consultation with the users. Thus both the predisposition of the techniques used in systems, and the implemented details were in harmony with their customer's vision of the ideal audience.

A very different example is supplied by CATO (Aleven 1997). Here the idea was to produce a system to support the teaching of case based argument. The system was evaluated according to improvements in student performance, and its own answers to coursework questions. The ideal audience here is very specific: the person grading the papers. Most academics have a clear idea of the structure they expect, the cases they want cited, and the points they want made about them. Thus building a system directed towards this audience is well placed to produce what is wanted, as CATO did. This is in no way intended to be a criticism of CATO, which is excellent system which fulfils its aims. On the other hand it is possible that the reason why this system has not had wider acceptance among the US legal community may be that its target ideal audience was too specific. There may well be a difference between the argument of a bright student which will attract a high grade, and the argument that an experienced attorney will produce.

A third example is given by the Split-Up system (Stranieri and Zeleznikow 1999). This system relies heavily on the use of an artificial neural net, which one might doubt would gain ready acceptance in the legal community. But the application is a rather special one: apportioning resources following a family breakdown. In this domain where the ideal audience seems to require *factor balancing* as the central reasoning technique, ANNs seem a natural technique to use. Again it was developed in very close co-operation with its potential users.

These three systems use very different techniques, but in each case the technique seems to have considerable affinity to the ideal audience that will receive its output. Moreover, and perhaps more importantly, the systems were developed from within the legal culture they were targeted at: there was therefore no mis-match between the construct of the ideal audience of the system builder and the user. I think there are some clear lessons here for those who wish to develop systems for practical use.

6. The Ideal Audience and the Practice of AI and Law

So far I have used the notion of the ideal audience to explain differences in approach amongst workers in AI and Law, and why some systems seem to have proved more acceptable to the target users than others. I now want to consider how the notion of an ideal audience can throw light on some problems that arise for researchers in an inter-disciplinary field like AI and Law.

The major output of academic research is research publications, papers in journals and conferences and books. These are written for an audience comprising largely academic research peers. Such an audience is very different from that of the target users of systems. In an inter-disciplinary field such as AI and Law, moreover, we must be aware that that the ideal audience for a computer science paper is very different from the ideal audience for a law paper, (different footnoting conventions are a superficial but real sign of this), and that the interdisciplinary audience will contain features of both. The danger here is that work may be directed too much at one sector of the audience. For example, papers intended for an AI and Law audience may address computer science concerns, with insufficient acknowledgement of the legal audience. The famous British Nationality Act paper (Sergot et al 1996) was addressed first and foremost at a logic programming audience. The hostile reception of some lawyers, who pointed to the simplicity of its "black letter" legal reasoning and found it incapable of addressing the subtleties of legal reasoning should have been expected, and perhaps addressed at the outset. The problem is that attention can be diverted from the promise of the technique because it is presented in a way which lays it open to attacks based on a misconstruction of its arguments. Influential as this paper has been, an impression that computer scientists working in this field try to impose their techniques without sufficient regard to legal concerns remains in some quarters (e.g. Moles 1992). Similarly papers on legal theory are often ignored by computer scientists when such papers make no effort to include them in the audience addressed: for example, any such proposals need to be aware of issues such as computationally feasibility. Such cases illustrate failure to take into account the range of the potential audience. Other examples can be seen where techniques interesting to computer science, such as objects, neural networks, agents, and the like, have been applied to AI and Law in ways interesting for the sake of the technique rather than for any application leverage. This techniques have, perhaps in consequence, attracted less attention than might have been expected in AI and Law.

There are, however, times when computer science concerns rightly take centre stage. If one intends to build a system to model legal reasoning, one must be aware not only of the ideal legal audience at which the output will be targeted, but also that computer science concerns are very relevant because one is *building a computer system*. Thus the kind of software engineering concerns that underlie much AI and Law work on the representation of source documents (e.g. Bench-Capon and Coenen 1992) are legitimate for software developers, even though they may not seem interesting to lawyers. Fidelity to legislative text is primarily to enable better systems to be built, not to allow better reasoning to be performed. Both writers and readers of work in AI and Law need to recognise this spread of concerns and the different constructs of the ideal audience that may be present in an inter-disciplinary field, and to which papers may sensibly be directed. Writers need to be aware of the different perspectives from which their work will be evaluated, and readers must be prepared to allow for the fact that work may address concerns that they do not share. Only in this way can the interchange of ideas be maximised.

Finally I want to make some brief remarks about how we might make the notion of an (ideal) audience explicit in AI and Law systems. Most of the dichotomies identified by Christie have been and will continue to be built into systems implicitly. There have, however, been some efforts, by myself and others, to make explicit the construction of the values and system of preferences embodied by the ideal audience as a way of deciding between competing arguments. This work has been explicitly motivated by the views of Perelman, and has drawn inspiration from a seminal paper of Berman and Hafner (1993), which drew attention to the need to look beyond the facts of precedent cases to the value preferences they revealed in order to resolve conflicting arguments. One such attempt is Bench-Capon and Sartor (2000). The idea here is that competing arguments are preferred according to the values of the audience to which they are addressed. These values, and their ordering, are revealed in previous decisions accepted by that audience. The idea is that reasoning about a case involves first constructing, from past decisions, an ordering on values, and a set of rules related to those values, which together form a theory, attributed to the ideal audience, which explains past cases and which should be applied in the current case. In the current context we can see this approach as explicitly recognising that the ideal audience is a construct which is built from the experience of the legal culture through the past decisions made within that culture. I believe this to be a very stimulating idea which opens new vistas for AI and Law.

7. Conclusion

The concept of an ideal audience is, as I hope I illustrated in the last section, a useful analytical tool for thinking about systems in AI and Law, and the practice of AI and Law. Arguments are directed towards an audience, and we need to characterise an ideal audience if we are to escape the traps of regarding arguments as subjective, or deriving their authority only from the fact of widespread acceptance.

References

Aleven, V., (1997). *Teaching Case-Based Argumentation Through a Model and Examples*. PhD Thesis, The University of Pittsburgh.

Bench-Capon, T.J.M., and Sergot, M.J., (1988). Towards a Rule Based Representation of Open Texture in Law, in C. Walter (ed), *Computer Power and Legal Language*, Quorum Books: Westport Conneticut, pp 39-60.

Bench-Capon, T.J.M., (1997) Argument in Artificial Intelligence and Law, *Artificial Intelligence and Law*, Vol 5 No 4, pp249-61.

Bench-Capon, T.J.M., and Coenen, F.P., (1992). Isomorphism and Legal Knowledge Based Systems, *Artificial Intelligence and Law*, Vol 1 No 1, pp 65-86.

Bench-Capon, T.J.M., and Sartor, G., (2000). Using Values and Theories to Resolve Diagreement in Law. In Breuker, J., Leenes, R., and Winkels, R., (eds), *Proceedings of JURIX 2000*. IOS Press: Amsterdam.

Bench-Capon, T.J.M., (2001). Truth and Consequence: Complementing Logic with Values in Legal Reasoning. *Information and Communications Technology Law*, Vol 10 No 1, 222001.

de Bakker, K,F,C., and Wassink, J.G.J., (1991). Development, implementation and impact of the TESSEC expert system. In *Proceedings of the European Group of Public Administration 1991*.

Berman, D.H., and Hafner, C.D., (1993). Representing Teleological Structure in Case Based Reasoning. In *Proceedings of the Fourth International Conference on AI and Law*, ACM Press: New York. pp 50-59.

Moles, RN., (1992). Expert Systems - The Need for Theory. In *Proceedings of JURIX 1992.* Koninklijke Vermade: Lelystad, pp113-122.

Perelman, Ch., and Olbrechts-Tyteca, L., (1969). *The New Rhetoric.* University of Notre Dame Press: Notre Dame.

Prakken, H., and Sartor, G., (1998). Modelling Reasons with Precedents in a Formal Dialogue Game. *Artificial Intelligence and Law*, Vol 6 Nos 2-4, pp231-87.

Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P, and Cory, T., (1986). The British Nationality Act as a Logic Program. *Communications of the ACM* 29, 5 pp 370-386.

Stranieri, A., and Zeleznikow, J., (1999). The Evaluation of Legal Knowledge Based Systems. In Proceedings of the Seventh International Conference on AI and Law. ACM Press: New York. pp 18-24.