

Thirty Years of *Artificial Intelligence and Law*: The First Decade

Guido Governatori · Trevor Bench-Capon ·
Bart Verheij · Michał Araszkiewicz · Enrico
Francesconi · Matthias Grabmair

the date of receipt and acceptance should be inserted later

Abstract The first issue of *Artificial Intelligence and Law* journal was published in 1992. This paper provides commentaries on landmark papers from the first decade of that journal. The topics discussed include reasoning with cases, argumentation, normative reasoning, dialogue, representing legal knowledge and neural networks.

Keywords Reasoning with cases, deontic concepts, knowledge representation, dialogue, neural networks

1 Introduction

Artificial Intelligence and Law began to emerge as a separate field of study following the first International Conference on Artificial Intelligence and Law (ICAAIL) held in Boston in 1987. As it developed through the second and third ICAILs, a number of topics established themselves as central to the field. Three of these were represented by papers in the very first issue of *AI and Law* journal in 1992 and all three of these papers are discussed in this paper.

Reasoning with legal cases is represented by Skalak and Rissland (1992), commented on by Trevor Bench-Capon. Reasoning with legal cases had been introduced in ICAIL 1987 by Rissland and Ashley (1987) which described the HYPO

Guido Governatori
Brisbane, Australia

Trevor Bench-Capon
University of Liverpool, UK

Bart Verheij
University of Groningen, The Netherlands

Michał Araszkiewicz
Jagiellonian University, Kraków, Poland

Enrico Francesconi
National Research Council of Italy (IGSG-CNR, Florence) and European Parliament (Luxembourg)

Matthias Grabmair
Technical University of Munich, Germany

system. While Ashley went on to develop the CATO system with Vincent Aleven (Aleven and Ashley, 1994), Rissland developed CABARET with David Skalak. A key feature of CABARET is that it is a hybrid system in which the reasoning with cases is deployed within a framework of rules. Reasoning with legal cases continued to be pursued throughout the decade and is also the subject of Hage et al. (1993), commented on by Bart Verheij, and Prakken and Sartor (1998), commented on by Trevor Bench-Capon. Hage et al gives a characterisation of what makes a legal case *hard*, a notion introduced in Gardner (1987). Prakken and Sartor provide a way of representing precedent cases as sets of rules, laying the foundations for much subsequent investigation of reasoning with precedents¹.

Modelling deontic concepts is the topic of Jones and Sergot (1992), commented on by Guido Governatori. In particular they identify where deontic modelling is essential, namely where the possibility of violation needs to be considered and reasoned about. Deontic concepts are also analysed in Sartor (1992), again commented on by Guido Governatori, where the focus is on normative conflicts.

The third topic represented in the first issue of the journal is the representation of legal knowledge (Bench-Capon and Coenen, 1992), commented on by Michał Araszkievicz and Enrico Francesconi. The particular concern of this paper was that such knowledge had to be maintained to cope with legislative change and it argues that this could be facilitated by maintaining correspondences between the source and the represented knowledge. Legal knowledge representation became more prominent in the second decade with the growth of interest in ontologies².

Another topic which was important in the first decade was the use of dialogues to model legal procedures and legal argumentation. Gordon (1993), commented on by Guido Governatori, which played an important part in popularising the use of dialogues in AI and Law, used the technique to model a particular legal procedure. Dialogues are also central to both Hage et al. (1993) and Prakken and Sartor (1998).

Dialogues are often motivated by the need to model argumentation. Argumentation is prominent in Skalak and Rissland (1992) and is the subject of Loui and Norman (1995), discussed by Bart Verheij, which shows how it can be useful to expand arguments to uncover hidden steps. During the second decade, modelling argumentation became dominated by the use of argument schemes³, and explicit representation of dialogue declined in importance.

Throughout the decade there was also interest in sub-symbolic techniques, especially neural networks. This interest is represented by Stranieri et al. (1999), commented on by Matthias Grabmair. This paper represents the culmination of a sustained body of work, and is notable for its recognition of the need to explain the recommendations, and the use of argumentation to do this. As discussed in Villata et al. (2022), elsewhere in this issue, sub-symbolic techniques have now become widespread in AI and Law, and the need for explanation remains a pressing concern.

The papers discussed here represent significant contributions to a range of topics that emerged during the formative years of AI and Law.

¹ See for example Section 10 of Sartor et al. (2022) elsewhere in this issue.

² See sections 4 and 5 of Sartor et al. (2022) and section 2 of Araszkievicz et al. (2022) elsewhere in this issue.

³ See section 3 of Sartor et al. (2022) elsewhere in this volume.

2 Arguments and cases: An inevitable intertwining (Skalak and Rissland, 1992). Commentary by Trevor Bench-Capon

When the AI and Law community came together at the first ICAIL in 1987, it was perceived that there was a division between approaches based on the formalisation of legislation in rule-based languages such as Prolog (e.g. Bench-Capon et al. (1987), Sherman (1987) and Biagioli et al. (1987)), and those concerned with the modelling of precedent cases (e.g. Rissland and Ashley (1987) and Goldman et al. (1987)). Often this division was seen as reflecting the difference between the Civil Law system prevalent in Europe and the Common Law system of the USA. The rules versus cases debate was the topic of a panel both at the 1987 ICAIL and the 1991 ICAIL (Berman, 1991). One major contribution of Skalak and Rissland (1992), which describes arguments and their theory as embodied in the CABARET system (Rissland and Skalak, 1991), was to show that this was a false dichotomy: not only are the two approaches not in conflict, but they are complimentary to one another, and the use of both in combination is *inevitable* in a full account of legal argumentation.

Edwina Rissland had previously developed her ideas about reasoning with legal cases with her PhD student, Kevin Ashley, on the HYPO project (Rissland et al. (1984) and Rissland and Ashley (1987)). HYPO addressed the domain of US Trade Secrets Law. While Ashley went on to develop CATO (Ashley and Alevan (1991) and Alevan and Ashley (1994)), also in the Trade Secrets domain, with his PhD student, Vincent Alevan, Rissland continued her exploration of reasoning with legal cases in the CABARET project with her new PhD student, David Skalak, and this work forms the basis of Skalak and Rissland (1992). For CABARET, the domain was changed to Home Office Deduction, which concerns a Federal income tax allowance for people who need to provide themselves with office facilities at home.

Unlike Trade Secrets, Home Office Deduction is governed by a specific piece of legislation, namely Section 280A of the Internal Revenue Code, and the paper focused on Section 280A(c)(1) which contained the essence of the statute. Section 280A(c)(1) comprises three clauses and seems to lend itself to a straightforward rule-based representation, in the manner of Sergot et al. (1986). Such rules, however, contain terms which require interpretation: for example the home office must be used *exclusively* as the *principal place of business* for the *convenience of the employer*. All the italicised terms must be established through case law: drafters of legislation typically leave room for interpretation in the light of the circumstances of particular cases, since it is rarely possible to define conditions which will give the desired outcome in every possible situation in advance of the consideration of particular states of affairs. This means that there will be situations not clearly covered by the rules: in the famous phrase of Gardner (1987), “the rules run out”. Thus while the reasoning must be to some extent what they term *specification-driven* to identify *what* claimants must establish to make good their entitlement, e.g. that the office is used for the convenience of the employer, it must also be *case-driven* in order to establish *how* the particular facts of the situation meets these requirements, and this should be consistent with decisions in previous cases.

“By ‘specification-driven’ we mean that the arguer specifies what the ideal case for his situation might be, independently of what cases exist in the

case base. By ‘case-driven’ we mean that the arguer creates and defines arguments according to what cases actually exist and are available to him.” Skalak and Rissland (1992), p 10.

In pure rule-based approaches such as Sergot et al. (1986), the case-based element was left entirely to the user. In HYPO, the aim was to generate arguments which the user would assess, and so the specification was in the mind of the user. But for a complete system, both are needed. Ashley and Alevan (1991) also recognised in CATO that the structure of the domain required representation and organised its factors into a hierarchy with the base level factors children of more abstract factors, above which are *issues*. Issues play the role of the terms needing interpretation found in legislation in CABARET - the things which must be established in order to prove a case. Although CATO organised its arguments by issue, the rule-based component was not explicit because, like HYPO, the aim was not to decide a case, but only to present the arguments. The framework of rules in which the factor based reasoning is deployed was only made explicit for CATO when the concern became prediction in Issue Based Prediction (IBP) (Brüninghaus and Ashley, 2003). IBP formed a significant component of the system described in Ashley and Brüninghaus (2009)⁴.

This structure for domain knowledge, with a top layer of rules interpreted through, and interleaved with, reasoning with cases, is now the acknowledged way of representing legal knowledge. For example it was the basis of Grabmair’s value judgement formalism (Grabmair, 2016) and methodologies such as ANGELIC (Al-Abdulkarim et al., 2016).

As well as establishing the structure that would become standard for representing legal domains, high level rules taken from statutes or *framework* precedents (Rigoni, 2015) interpreted using precedent cases, Skalak and Rissland (1992) was also influential in its emphasis on argumentation, an aspect which has become increasing central to AI and Law (and AI in general) over the last three decades. The emphasis in this paper, however, is, like HYPO, very much on the *generation* of arguments rather than their assessment to predict and justify outcomes, which was the typical emphasis in later approaches.

Arguments are generated in Skalak and Rissland (1992) with a three-tiered approach using *argument strategies*, realised using *argument moves*, which are implemented using *argument primitives*. The appropriate strategy is selected by reference to the rule governing the case and the point of view. If the rule conditions are met and the point of view is positive, the hit must be confirmed but if the point of view is negative, the rule must be discredited. If the rule conditions are not met, the miss must be confirmed for a negative point of view, or the rule broadened for a positive point of view. The move is determined by the precedents available and their dispositions.

Once the strategy has been selected, the precedents are used to select a move. Depending on the outcome in the precedent and the strategy being employed the precedent must be analogised to or distinguished from the current case. These moves are then implemented through detailed comparison of the features of the current case and the precedent to determine the degree and nature of the matches

⁴ See Section 7 of Sartor et al. (2022), elsewhere is this issue.

and mismatches between the two⁵. For instance, when broadening a rule, citing a precedent with the desired outcome that also failed to satisfy a rule antecedent, can be used to argue that since the missed condition was not necessary in that case, it is not needed in this case either.

In Skalak and Rissland (1992), the argument moves are limited to those which can be produced using the form of argument the authors term a *straightforward argument*, in which the facts of a current case are compared with a precedent case with the desired outcome. The paper, however, gives a taxonomy of argument forms used in legal argumentation, which include a variety of additional forms of argument.

Some of these other forms have also been widely studied: notably *slippery slope* arguments (e.g. Walton (2015)), and *Weighing or balancing arguments*. The latter has been extensively studied in AI and Law both qualitatively (e.g. Lauritsen (2015) and Gordon and Walton (2016)) and qualitatively (e.g. using Reason Based Logic (Hage et al. (1993)⁶ and Hage (1996)), Bench-Capon and Sartor (2003) and Bench-Capon and Prakken (2010)).

Some of the other forms have not, however, been much studied outside Skalak and Rissland (1992) itself. These argument forms are mainly rhetorical. *Straw Man* is used disparagingly by an opponent to describe an argument with an allegedly obvious flaw. *Make weight* and *'throw the dog a bone'* are similarly weak and can only be used in conjunction with stronger arguments. *Hedging* arguments are only used in case the main argument fails. Finally there is the *Turkey, chicken and fish or 'double negative' argument*, where it is argued that the current case is so unlike the cases where a rule's conditions were held not to have been established that the rule should apply to the current case. Later argumentation approaches such as Prakken and Sartor (1998)⁷, tend to be normative, defining a rational procedure for testing the tenability of a claim, rather than generating realistic disputes, and so do not consider such rhetorical aspects. The paper concludes with a detailed discussion of how these different forms can be deployed in the CABARET framework to enrich the argumentation.

Rissland and Skalak further pursued (with M Timur Friedman) argument generation in the BankXX project (Rissland et al. (1996) and Rissland et al. (1997)). This interesting project used heuristic search to construct arguments by finding building blocks, such as collections of facts, cases as dimensionally-analyzed fact situations, cases as bundles of citations, and cases as prototypical factual scripts, as well as legal theories represented in terms of domain dimensions. BankXX did not have many successors, but the argument style proposed in Skalak and Rissland (1992) remains current. The notion of producing an argument with a hybrid approach, using a rule-driven strategy implemented using case-driven argument moves, remains a central way of providing justification of case outcomes. In the 90s this was mainly pursued using dialogue games. Examples include Hage et al. (1993)⁸, Prakken and Sartor (1996), Prakken and Sartor (1998)⁹, and Bench-Capon et al. (2000). Later, however, as computational models of argumentation

⁵ These primitives play the role of the factor partitions in Wyner and Bench-Capon (2007) and the functions in Prakken et al. (2015) used in the instantiation of their argument schemes.

⁶ See Section 7.

⁷ See Section 9.

⁸ See Section 7.

⁹ See Section 9.

became prevalent, argument schemes (essentially the *forms* of Skalak and Rissland (1992)), as introduced in Walton (1996), became widely used. Here the scheme plays the role of the rule, with the critical questions indicating the moves to be made in response, but there is a wide variety within this basic approach. Examples include Verheij (2003b)¹⁰, the three approaches to representing *Popov v Hayashi* in a special issue of the journal (Atkinson, 2012)¹¹, (Bench-Capon (2012), Gordon and Walton (2012), Prakken (2012)), Prakken et al. (2015) and Grabmair (2017), all of which reconstruct the arguing with cases in terms of argument schemes. This form of argumentation is now also advocated as a means of explaining the prediction of Machine Learning approaches in e.g. Branting et al. (2021)¹².

The current use of argumentation in AI and Law focuses on the inferential rather than the rhetorical aspects and so several of the forms identified by Skalak and Rissland are no longer considered. None the less the influence of the paper, both in terms of bringing together rule and case based approaches to provide a coherent structure for knowledge of legal domains, and its emphasis on argumentation, studied through an analysis of the different forms/schemes used to generate arguments, remains very significant.

3 Deontic logic in the representation of law: Towards a methodology (Jones and Sergot, 1992). *Commentary by Guido Governatori*

Although this paper was published in the first volume of the Journal, its topic and contribution are still very relevant for the development of AI and Law solutions. Jones and Sergot discuss the role of deontic logic for the representation of norms. The debate repeatedly recurs with new systems and implementations which ignore the discussion being proposed, including recent approaches from the current Rules as Code movement (Shein, 2021), whose aim is exactly the representation of law as computer programs. Jones and Sergot (1992) focuses exactly on this point and argues that, in general, deontic logic is required to succeed in this task. Nevertheless, and correctly, the contribution does not advocate any specific logic.

Deontic Logic was proposed by von Wright (1951) as a logic to model normative expressions on obligations, permissions and related notions. It received substantial attention from the philosophical and computer science communities (with a conference dedicated to it, DEON, e.g. Liu et al. (2021)). However, as the paper mentions, no applications had been created at the time the paper was written, and the investigation concerned how deontic logic(s) could be applied. The situation has not changed: studies on how to use it for legal reasoning have proliferated, but with some notable exceptions, applications have not been developed¹³.

The argument in favour of deontic logic in the paper runs as follows: First, the paper analyses the claim (Bench-Capon, 1989) that modal deontic expressions

¹⁰ See section 3 of Sartor et al. (2022) elsewhere in this issue.

¹¹ See Section 9 of Sartor et al. (2022), elsewhere in this issue.

¹² See Section 7 of Villata et al. (2022), elsewhere in this issue.

¹³ Since 1992, the journal has published 19 articles discussing techniques based on or adopting some form of deontic logic. Of those contributions, only four of them discuss actual applications based on deontic logic (contracts (Governatori et al., 2018), decision support (Islam and Governatori, 2018), and compliance management (Hashmi and Governatori (2018), and Ardila et al. (2021)).

(e.g., “must”, “shall”, “ought”) often do not necessarily require to be represented by deontic modal operators. Moreover, they agree that many fragments of law can be successfully modelled without deontic logic and deontic operators for practical purposes. Then, they mention the work by Herrestad (1991) to point out that there are situations that cannot be properly represented without deontic logic. Then the paper investigates the cases requiring deontic notions: Jones and Sergot agree that, in general, it is possible to distinguish two types of norms: definitional or constitutive norms (called qualification norms in the paper) and regulative norms, and that for specific situations and the purpose of a particular application “a decision has to be made about what needs to be represented and what not” (Jones and Sergot, 1992, p. 55). Thus, for some applications, a purely definitional formalisation might be adequate. Then, they delineate what type of situations benefit from the explicit formalisation of deontic notions.

The situations where deontic notions are beneficial are identified as cases where the provisions account for violations and when violations of (primary) obligations trigger other obligations. Thus, the paper argues about the *ideal vs actual* dichotomy: obligations (and other deontic notions) describe what is ideal (according to a normative system) and it is not necessarily accomplished in actual cases. Hence, in general, what is actual and what is ideal do not necessarily coincide. However, those two aspects must co-exist and determine further legal consequences; typically, normative systems contain provisions to account for penalties and other compensatory measures (called Contrary-to-Duty obligations). To illustrate these notions, the paper addresses two possible application lenses where regulations (normative systems) can be used (1) as system specifications or (2) as sets of norms directed to users. The two aspects are illustrated by the same (simple) example (part of the Imperial College Library regulations). While the purposes of the two scenarios are different, they share the same concern: how to address situations diverging from the ideal case (though for different reasons). Hence, the paper concludes that deontic operators are advantageous when building a system that conforms with norms or which provides norms for the user, used by them to decide on how to act.¹⁴

Finally, the paper identifies some methodological aspects to determine whether deontic logic is useful for a formalisation. The first consideration is that while often deontic concepts are introduced by expressions with a modal auxiliary such as ‘must’, ‘should’, ‘ought’, the linguistic expression is just a surface indicator and “*never a faithful guide for content*” (Jones and Sergot, 1992, p.62). Such expressions can be used with a different meaning, and expressions without the auxiliary can convey the same meaning (for instance, most provisions in the Italian Criminal Code adopt the expression “whoever does ... is punished with ...” to express prohibitions, and modern Australian legislative instruments typically use the form “it is an offence under Section X to do Y when Z”). Thus, for Jones and Sergot, the main considerations to determine whether to use deontic operators are:

1. Is there any meaningful sense in which the provision at hand can be said to be violated?

¹⁴ For example, questions of conforming to norms have arisen recently in the context of autonomous vehicles, e.g. (Prakken, 2017; Bench-Capon and Modgil, 2017; Bhuiyan et al., 2020)) and contracts (Governatori, 2005).

2. Does such violation result in a new state of affairs which needs to be represented explicitly for the purposes at hand? (Jones and Sergot, 1992, p. 59)

The contribution briefly touches on the issue of using a rudimentary version of deontic logic where the deontic expressions are conflated in, for example, propositions allowing “the consistent assertion of ‘oughtA’ and ‘not-A’ and the derivation of ‘permittedA’ from ‘oughtA’, but which otherwise remained silent on the logical properties of the deontic modalities” (Jones and Sergot, 1992, p. 57). They argue that the purpose of the formalisation should dictate this, but the risk is to miss some challenging problems that could emerge. But, when one engages in deontic logic, non-trivial questions have to be addressed (the paper mentions the choice of what detachment principles to use and whether they have to be taken in a restricted form or not (Jones and Sergot, 1992, p. 58)).

As we alluded to above, the paper mentions the result by Herrestad, where adopting a first-order logic for the representation of the deontic concepts could lead to counterintuitive results. What about different logics? Governatori (2015) shows an example where a naive and apparently intuitive formalisation of a legal scenario in (Linear) Temporal Logic ends in some paradoxical results. However, it is important to notice that the findings reported in Herrestad (1991) and Governatori (2015) are not impossibility results. Indeed, Governatori (2015) discusses a procedure to recover a meaningful representation in (Linear) Temporal Logic. Knowing the expected legal outcome of a scenario allows us to provide suitable formal representation in a chosen logic. However, this is somehow against the aim of the formalisation itself, since the representation is to be used to predict and investigate the possible outcome of a legal case.

To conclude, the lesson learnt from this seminal contribution by Jones and Sergot is that deontic logic is essential for the full understanding and representation of normative systems. However, specific applications can dispense with it when their purpose is somehow limited and can be restricted to the qualification aspects of normative systems. Furthermore, a similar argument applies to non-deontic logic. Still, suitable formalisations in such logics might fail to provide appropriate conceptual models of the normative system that is represented or not provide sound results especially when issues of conforming to or violating the norms arise. Finally, careful consideration should be paid to the properties of the deontic operators, and thus, their study, after thirty years of this contribution, is still relevant and still needed.

4 Isomorphism and legal knowledge based systems (Bench-Capon and Coenen, 1992). *Commentary by Enrico Francesconi and Michał Araszkiewicz*

Bench-Capon and Coenen (1992)¹⁵ which appeared in first issue of *AI and Law* journal, represents a very important reference point in the AI and Law domain, in particular in the field of legal knowledge representation. It was an effective contribution to the early '90s debate about the extent to which legal reasoning can be reduced to reasoning with rules.

¹⁵ This was an extended version of Coenen and Bench-Capon (1991).

According to Thorne McCarty, a possible solution to such a question is the development of systems based on “deep conceptual models” of the relevant domain McCarty (1984). On the other side, Bench-Capon and Coenen (1992) argued that, for most practical applications, an expert system can be based on a formalisation of the legislation itself, which should be a *faithful* representation of the legal rules. This introduced the key concept of *Isomorphism*, which inspired much AI and Law research in the following years. Far from being opposite, these positions can be seen as complementary, because they look at the same objectives from two perspectives: top-down and bottom-up, respectively.

Bench-Capon and Coenen (1992) in particular introduces and contributes to the identification the characteristics of the concept of “isomorphism”, as an approach to building knowledge-based systems in the legal domain, starting from a faithful representation of data, according to a well-defined correspondence between source documents and the representation of the information they contain used in the systems.

According to Bench-Capon and Coenen, one of the main benefits of the implementation of “isomorphism” is the separation between the data and the logic managing such data. In particular, following an isomorphic approach to create a legal knowledge base and a related set of rules, it is possible to say of any item in the knowledge base that it derives from some self-contained unit in the source material. A main motivation is to facilitate maintenance, which is highly desirable given the frequency of amendments to legislation.

This idea gave rise to an interesting debate on this matter, as witnessed by the objection of Marek Sergot (based on the experience of Sergot et al. (1991) and reported in Bench-Capon and Coenen (1992)). Sergot argued that it is appropriate to follow an isomorphic approach if the legislation is itself well structured, but otherwise this approach might become cumbersome. However, legislation is very often not well structured¹⁶. In such a case, isomorphism would lead to a poorly structured knowledge base, one which fails to correspond to the ‘real world’ problem.

Nevertheless, Bench-Capon and Coenen considered an isomorphic approach to be valid, because many of the problems with expert systems come from a mismatch between the rule based conceptualisation of the expert systems and the conceptualisation of the user. The structure of legislation is what we have to manage and it is this that is familiar to the users; moreover, all legislation is typically subject to modifications and having a knowledge base with a different structure will complicate maintenance. According to the authors, isomorphism has advantages in terms of the development methodology, verification, validation, use and maintenance of systems, as it gives a correspondence between user and system conceptualisations, and so makes the behaviour of the system more transparent to the user.

In the paper a practical case study of isomorphic development is presented related to the development of a regulation based Knowledge Based System (KBS) for British Coal using the Make Authoring and Development Environment (MADE) (Coenen and Bench-Capon, 1992).

¹⁶ For example, The Indian Pension Rules modelled in Sergot et al. (1991): “In common with many other examples of legislation and regulations, especially those that refer to periods of time, the Pension Rules are imprecise and very casual about many of the key concepts. They are certainly not precise enough to be formulated directly as an executable program” (p 119).

The result of the analysis of sources of law is a set of rules (Rule Base) and a hierarchical set of objects (Class Hierarchy). Such a knowledge base is codified in an intermediate representation MIR (Make Intermediate Representation), based on Bench-Capon and Forder (1991), which is amenable for computation, while allowing the rule definitions to be traced back to the original source material and vice-versa.

It is worth noting that the intermediate representation (MIR) involving a set of objects (Class Hierarchy) and related properties, assumes a form that would be later known as an “ontology”. Moreover, MIR is compiled in an executable target representation (CMIR) allowing reasoning, as do current reasoners for ontologies (e.g. Pellet (Sirin et al., 2007)). Therefore, Bench-Capon and Coenen (1992) describes a knowledge representation approach which will be later considered an inspiration of other works using semantic web technologies for knowledge modeling and tools for extracting knowledge from legal texts¹⁷.

The isomorphic approach to the development of legal knowledge bases has been acknowledged in influential contributions to the state of the art in the 1990s, including Gordon’s monograph on computational modeling on procedural justice (Gordon, 1995) and Prakken’s volume on logical tools for modeling of legal argument (Prakken, 1997a); in the latter case, the term “structural resemblance” was preferred but it was also indicated that it covers all the various senses of the term “isomorphism” used in the earlier literature. Although focus switched to ontologies from the mid 90s, isomorphism was revisited in Bench-Capon and Gordon (2009) where it was shown that several different representations of a given piece of legislation could be considered isomorphic and that such differences matter, in terms of where the burden of proof is allocated, in terms of the explanations produced, and in terms of the operational procedures that are reflected.

5 Normative conflicts in legal reasoning (Sartor, 1992). *Commentary by Guido Governatori*

Sartor (1992) addresses the issue of how to use logic and formal systems to handle normative conflicts in the context of legal reasoning. The first topic examined in the paper is the type and sources of normative conflicts. Specifically, three causes of “incompatible legal qualifications” are identified (Sartor (1992), p209):

1. the dynamics of the legal systems,
2. the legal protections of conflicting interests, and
3. the uncertainty concerning the content of legal sources.

A key issue is that “not all normative conflicts can be [...] or should (given the positive function of certain normative conflicts, such as those between rules and exceptions) be prevented by the legislator” (Sartor (1992), p209).

It then proceeds with a methodology to map textual provisions/norms into a formalisation.

¹⁷ The MIR methodology was followed by Pepijn Visser in developing his own ontology (Visser, 1995) for Dutch Unemployment law. Subsequently Visser and Bench-Capon worked together on a series of papers on ontologies including Visser and Bench-Capon (1998) which is discussed in section 2 of Araszkievicz et al. (2022), elsewhere in this volume.

$$\text{name}(x_1, \dots, x_n): \text{conclusion} \leftarrow \text{condition} \quad (1)$$

that can be represented by

$$\text{name}(x_1, \dots, x_n) \rightarrow (\text{conclusion} \leftarrow \text{condition}) \quad (2)$$

where x_1, \dots, x_n are the variables appearing in *conclusion* and *condition*, \rightarrow is the classical material implication of first-order logic, and \leftarrow is the logic programming implication. In addition, it is possible to specify an ordering among the rules. Thus, the expression $\text{name}_1 > \text{name}_2$, where name_1 and name_2 are the predicates naming two rules, means that rule name_1 takes precedence over rule name_2 in case they both apply and the conclusions conflict.

The paper argues that the representation and the naming convention are particularly useful to capture exceptions and identifies two types of exceptions: *exceptions to norms* and *exceptions to effects*. We have an exception to a norm when the norm is unambiguously identified, but it does not apply in a particular situation. An exception to an effect is when a situation does not entail the legal qualification (conclusion) of a norm (Sartor (1992), p213). An advantage of this formulation, which has not been widely adopted in subsequent work on the formalisation of normative systems, is that it does not require different forms to distinguish *rebutting defeaters* (exceptions to effects) from *undercutting defeaters* (exceptions to norms). A rule for $\neg\text{conclusion}$ is a rebutting defeater of (1), while a rule for $\neg\text{name}$ is an undercutter defeater of the rule in (1).

After setting up the theoretical background, the paper presents in Section 2 three case studies (based on the Italian legal system) to ground the discussion on the conflicts and show how to address the issues using the proposed formalisation.

The first example concerns the dynamic aspects of legal systems, where the paper argues that a way in which a legislator can “eliminate a certain legal content p in circumstances q ” is to enact the rule

$$n: \neg p \leftarrow q \quad (3)$$

and establish that the new rule is stronger than any rule for p (Sartor (1992), p215). Sartor further argues that this approach provides a minimal amount of modification to the legal system since only one new rule has to be introduced, and conflicts are avoided by the conflict resolution mechanism of the underlying formalism.

In the second example, Sartor discusses the issue of reasoning with rules and exceptions. Similarly to the previous scenario, the analysis is based on rules modelling (part of) Italian Tort Law. After the norms have been formalised based on the approach he proposed, the focus is on the issue of whether it is possible to represent exceptions using completions. For instance, given the rules

$$n(x): p \leftarrow q_1 \wedge \dots \wedge q_n \quad (4)$$

$$m_1(x): \neg p \leftarrow r_1 \quad (5)$$

⋮

$$m_l(x): \neg p \leftarrow r_l \quad (6)$$

where the ordering of the rule is $n > m_i$ (rule n prevails over rule m_i). It is possible to rewrite the rules as

$$p \leftarrow q_1 \wedge \dots \wedge q_n \wedge \neg r_1 \wedge \dots \wedge \neg r_l \quad (7)$$

with an (apparent) equivalent representation. Sartor argues that the two representations are not the same in the context of legal reasoning. He claims that the reformulation above is based on “the (erroneous) opinion that legal norms are only *prima facie* in conflict so that contradictions can be eliminated by means of interpretation” (Sartor (1992), p218). The consequence of this opinion is that we can translate any set of conflicting norms into a set of non-conflicting norms, and the two sets are logically equivalent for any case. He argues that this is not the case and gives two reasons:

- Non-monotonic reasoning;
- The dialectic of rules and exceptions.

For non-monotonic reasoning, the argument is that a legal-decision maker has to use the available information when a decision is made. However, if more details were available, the outcome could have been different. The two representations are equivalent provided the legal-decision maker has complete knowledge of all (possible) facts of a case; thus, the equivalence holds only for a logically omniscient legal decision-maker (for a discussion of this aspect see Governatori et al. (2009)).

For the dialectic aspect, the key issue is that the transformation is not resilient to modifications in the underlying legal system. Every time new rules are promulgated or norms are amended or abrogated, a new transformation must be created to ensure the correspondence between the provisions in the legal system and their formalisation.¹⁸

The third and final example illustrates a scenario where multiple interpretations are possible; each representation corresponds to an interpretation, where the legal decision-maker must first decide between the alternative interpretations before they can take a (justified) decision (Sartor (1992), p219).

The paper then presents two systems to deal with conflicting norms: Brewka preferred sub-theories (Brewka, 1991) and AGM theory change framework (Alchourrón et al., 1985). The two systems are designed for different purposes and intuitions. Brewka’s approach is to reason with a (possibly) inconsistent set of rules, and AGM to revise an inconsistent set of rules. For AGM, Sartor examines the notion of *safe contraction* proposed by Alchourrón and Makinson (1985) that introduced a (non-circular) ordering to determine the minimal set of elements (of a set of norms) for a conclusion p . Thus, if one wants to contract p (making p no longer a conclusion of the set of norms), safe contraction allows for the identification of the elements that can be “blamed” for p , and these elements can be discarded. Sartor studies conditions under which Brewka’s approach and AGM framework produce the same outcomes.

Then the paper briefly examines the logic proposed by Alchourrón (1986), based on the intuitions developed by Alchourrón and Makinson (1982), to handle

¹⁸ Incidentally, both Di Giusto and Governatori (1999) and Maranhão (2001) independently developed techniques corresponding essentially to the intuition outlined by Sartor (and extending it to cover the full propositional language) to compile conflicts away in the context of the revision of rules/norms when contradicting rules are present.

inconsistent rules and to reason with them to derive (a consistent set of conclusions), instead of revising the theory. Technically, the mechanism employed by the logic is the same as safe contraction. Both Brewka (1991) and Alchourrón (1986) use ordering over rules to solve contradictions. Sartor (1992) (p. 228) points out the differences between the two approaches and draws some conclusions on their use for legal reasoning. Brewka's approach is useful for *partial incompatibility*: "A norm n_1 is incompatible with stronger norms contained in the legal system, but there are possible cases where n_1 is applicable." Thus, partial incompatibility is suitable for *derogation sensu stricto*, all incompatible norms are valid, and their conflicts are handled by the hierarchical order of the legal system. Alchourrón's proposal corresponds to *total incompatibility*: when a norm is weaker than a conflicting norm there are no cases where the weaker norm can be applied; thus, the weaker norm is *tacitly abrogated*.

The final part focuses on how to establish the hierarchical ordering of a legal system. This follows from the analysis in the paper supporting the claim that "reasoning with inconsistencies in the law relies on the hierarchical ordering \leq over the legal systems" (Sartor (1992), p229). The paper shows how to map several well-understood criteria from legal theory (e.g., *lex superior*, *lex posterior*, ...) into an ordering over rules.

The analysis of the types of conflicts, the use of non-monotonic reasoning to represent exceptions, and the differences between derogation (use of exceptions) and abrogation (changes in the legal system) have provided important insights for the development of techniques for legal reasoning over the past thirty years. It is very relevant for the current investigation to properly understand what techniques are appropriate for specific legal reasoning applications.

6 The Pleadings Game (Gordon, 1993). *Commentary by Guido Governatori*

The Pleadings Game (Gordon, 1993) is the result of Gordon's doctoral research, and an extended version including all background material is available in his doctoral dissertation published as Gordon (1995). The work is described as an exercise in *computational dialectics* whose aim is to provide a computationally-oriented model of Robert Alexy's discourse theory of legal argumentation (Alexy, 1989).

The Pleadings Game is a game between two parties, the plaintiff and the defendant, meant to model the pleading part of a legal proceeding for a civil case (in the US legal system)¹⁹. However, Gordon remarks that his "model is more akin to the common law practice than the 'modern' law of civil procedure in the United States" (Gordon (1993), p241). The main difference is that in the latter the parties are not required to provide arguments, but only to assert or deny "essential facts" that they believe are relevant to the dispute. The aim of the game is to determine the legal and material facts in a case. The parties take turns in exchanging arguments about the issues of the case. First-order sentences and defeasible rules form the arguments; more specifically, the parties can concede, deny, defend claims, and declare defeasible rules. Thus, rules (defeasible or not)

¹⁹ To be precise, it comprises both the pleading and the trial, but in the trial, there is only one single player, the judge, and as remarked by Lodder in his review of Gordon's PhD thesis (Lodder, 2000), it is not a very interesting game.

can be asserted during the game. Accordingly, the parties can argue about the validity of the rules, and in this sense, the rules are pieces of “evidence” in the game.

The paper starts by highlighting the principles on which the game is based:

- No party can contradict itself;
- A party who concedes that a rule is valid must be prepared to apply it to every set of objects that satisfy its antecedent.
- An argument supporting an issue may be asserted only when the opponent has denied the issue.
- A party may deny any claim made by the opponent if it is not a necessary consequence of its own claims.
- A party may rebut a supporting argument for an issue it has denied.
- A party may defeat the rebuttal of a supporting argument for one of its own claims if the claim is an issue.

While, as we mentioned above, the model is based on Alexy’s framework (Alexy, 1989), it does not fully implement it: the differences are discussed on pages 242-3 of Gordon (1993). It is important to notice that the Pleadings game proposed by Gordon is not about *deciding* the issues for a case (this part is reserved for the trial game), but to *identify* the issues to be decided in the trial.

The issues to be decided are encoded as first-order formulas, and how these formulas interact is described by a set of rules corresponding to the encoding of Article Nine of the Uniform Commercial Code (UCC) of the United States. The formulas and the rules are used to form arguments and counterarguments. The computational model to determine the consequences of the arguments and counterarguments (and thus resolve the issue) is based on the non-monotonic logic of *conditional entailment* by Geffner and Pearl (1992). Gordon’s motivation for choosing such a logic is the ability to represent exceptions. Furthermore, he argues that Article Nine of the UCC requires both *explicit* and *implicit* exceptions, even though he notices that from a practical point of view, the vast majority are covered by explicit exceptions. Explicit exceptions are divided into two cases:

- a section does not apply because another section does apply;
- a section does not apply because specific conditions under which it does not apply are satisfied.

The logic of conditional entailment has at its heart the notion of a default, an expression of the form $p \Rightarrow q$, where p is the antecedent, and q is the consequent of the default, and it employs specificity to resolve conflicting defaults. However, for explicit exceptions, to “cancel” the applicability of a default instance δ under some condition q we can use the expression (if q then (not δ)). However, the paper provides notation for a more convenient encoding of the rules. Thus, for example, section 9-105 of Article Nine of the UCC, which defines what goods are, is encoded as follows

```
(rule s9-105 (x)
  if (movable x)
  then (goods x)
  unless (or (money x) (instrument x))).
```

where `s9-105` is the identifier of the rule, `x` is the parameter of the rule, the `if` part contains the conditions or antecedent of the rules, `then` specifies the conclusion of

the rules, and **unless** the exceptions to it. The formalism allows for the reification of rules; accordingly, it is possible to use the identifier of a rule or a rule instance as a term in another rule; a useful feature to facilitate the encoding of rules containing references to other rules²⁰. Furthermore, the paper shows the semantics of rules and their instances, how to model priority of rules, and how to encode legal principles such as *Lex Superior* and *Lex Posterior*.

After the detailed presentation of the computational model, the paper moves to the description of the procedure of the game, defining what moves (consisting of speech acts) the parties are allowed to play and termination conditions. To summarise, after a complaint is made, establishing the main claim (a proposition in the underlying logical language) of the game, the game begins with the first move by the defendant, where the defendant can concede or deny the claim. If the claim is denied, the plaintiff has to provide an argument supporting the claim. An argument is a set of formulas and a rule whose combination entails (or rejects) a proposition. If the plaintiff succeeds in making an argument, then the turn passes to the defendant, who can concede, deny, or defend an argument or a rule, or declare new rules. When a rule is conceded, the rule becomes valid and is no longer disputable. The game continues with alternate moves by the players until there are no relevant statements to be answered, at which point some issues may remain unresolved. Discussion of an issue ends when it has been (1) conceded, (2) denied twice, or (3) countered. When an issue is conceded or countered, then the two parties agree on it. When an issue has been denied, and the denial has been denied, it is left unresolved without further discussion at this stage, and it is passed to the trial.

The central part of the paper is dedicated to the detailed formalisation and presentation of a case based on Article Nine UCC. In addition to the formalisation of the piece of legislation, the paper describes the game corresponding to the case step-by-step.

The final part of the paper is dedicated to discussing a mechanism (and some related technical issues, e.g., computational complexity of the problem) to create (adduce) arguments and visualise them, including mechanisms to attach, rebut, and support arguments. Gordon himself went on develop these aspects in his *Carneades* system (Gordon, 2013). This final part can be seen as a contribution to the development of the field of formal argumentation, an example of AI research that largely started in the AI and Law realm²¹. Furthermore, the Pleadings Game was one of the first computationally-oriented models to represent a legal procedure as a dialogue game. It inspired a great deal of further work from a variety of researchers on the use of dialogue games to model different aspects of non-monotonic and legal reasoning (see, for example, Hage et al. (1993)²², Prakken and Sartor (1996), Prakken and Sartor (1998)²³, Bench-Capon et al. (2000), Hage (2000), Walton (2003), Prakken (2008), Wardeh et al. (2009) and Burgemeestre et al. (2011)).

²⁰ While this feature was proposed in other work contemporary to the work of Gordon, see, for example, (Sartor, 1992), a deeper investigation has been neglected until the recent work by Governatori and Olivieri (2021) that explicitly addresses the issue of modelling legal references.

²¹ For the contribution of AI and Law research to Argumentation and AI in general, see Bench-Capon and Dunne (2007).

²² See Section 7.

²³ See Section 9.

7 Hard cases: A procedural approach (Hage et al., 1993). *Commentary by Bart Verheij*

The paper *Hard Cases: A Procedural Approach* by Jaap Hage, Arno Lodder and Ronald Leenes appeared in the journal in 1993 (Hage et al., 1993). The paper has 55 pages and 11 sections (plus an appendix). It currently has been cited by 127 scholarly sources.²⁴

7.1 Overview

The paper's starting point is that the nature of legal reasoning is not logical, but procedural. Very briefly, in a logical view of legal reasoning, the outcome of a legal case is determined by given factual and legal input. In contrast, in the paper's procedural view of legal reasoning, the outcome of a legal case depends on what happens in the actual legal procedure and what the participants actually do. It is an assumption of the paper that the existence of hard cases shows the essentially procedural nature of legal reasoning.

The paper provides a theoretical discussion and defence of the procedural perspective on legal reasoning, followed by a characterization of a dialogical model of legal reasoning that can represent such a procedural perspective: Dialogical Reason Based Logic. A case study of an actual Dutch hard case about post-traumatic neurosis is used for illustration and assessment of the approach. Along the way, the topic of hard cases is treated from the perspective of legal theory, and then defined in terms of the paper's procedural approach. The paper further contains an extensive discussion of related work in the field of AI and Law by Berman and Hafner (1987), Gardner (1987) and Gordon (1991, 1993) (see also Hafner and Berman (2002) and Gordon (1995)). It is further discussed how the procedural approach to legal reasoning changes what is to be expected of a legal decision support system, and how such a system should be designed.

The paper is densely written and rich in ideas. A few are highlighted here: hard cases and good lawyers; procedural justice; and the design of legal decision support systems.

7.2 Hard cases and good lawyers

In the paper, the procedural view on legal decision making is applied to hard cases. Here a case is considered to be hard if determining the outcome of the case necessarily involves arational, non-logical reasoning. One can think of situations where judicial discretion is needed, e.g., when logically there is no outcome or there is an inconsistency, or of situations that are not logical in the classical sense, e.g., when the allocation of burden of proof, the validity of principles and rules, or the procedural rules themselves are at issue.

In other words, a case is hard if the outcome is not determined by logically valid reasoning on given factual and legal premises. To define hard cases, Dialogical Reason-Based Logic (DRBL) is presented, which is a framework of procedural

²⁴ Google Scholar, scholar.google.com, accessed February 18, 2022.

rules for two dialogue participants (representing the sides in a legal case) and an arbiter (representing a decision maker, such as a judge). The appendix lists 42 rules that semi-formally describe DRBL. A dialogue consists of dialogue moves by the participants and possibly one or more arbiter decisions.

It can happen that the participants in a debate do not need an arbiter, and can arrive at an outcome themselves. Then the case is easy or clear. But in a hard case, a debate participant can call for an arbiter decision, and if the other side does not object, the arbiter then provides a decision. The setting of DRBL hence allows for a definition of hard cases, as follows: A case is hard if, in the dialogue about the case, the arbiter actually decides about a claim (p. 135). The arbiter decision can concern the main topic of the dialogue, or a topic in a sub-discussion. In an example dialogue (p. 136), it is discussed whether the rule that loans must be repaid is excluded given the circumstance that the loan is more than thirty years old. In the example, a call to the arbiter is used to determine the validity of this exclusion rule.

This perspective on hard cases also gives a characterization of a good lawyer: someone who can show that a seemingly clear case is actually hard, and that a decision by an arbiter is required.

7.3 Procedural justice

The paper's approach is presented in the context of a Rawlsian discussion of procedural justice,²⁵ in which three categories are distinguished. The first is perfect procedural justice, where there exists an external standard of just outcomes, and a procedure that guarantees such an outcome. Dividing a cake is an example. The second is imperfect procedural justice, where there also exists an external standard of just outcomes, but procedures cannot guarantee such an outcome, and instead can only aim to approximate that. Fact finding in criminal law is an example. The third is pure procedural justice, where there does not exist an external standard of just outcomes, and instead the justness of an outcome is determined only by the correct application of the procedure. Elections are an example.

Hage et al. (1993) emphasizes that the approach presented is one of pure procedural justice. Quoting from the paper, 'Only by correctly applying the rules of law, we can arrive at valid legal conclusions, and these conclusions are valid because they are the result of correct application of the rules.' (p. 118) And: 'In our opinion, judging cases should be compared to elections and to gambling, rather than to dividing a cake or to a criminal procedure.' (on that same page).

The emphasis on pure procedure plays a key role in the discussion of the distinction with related work by Gordon (Section 9, p. 152-8 of the paper). For Gordon, hard cases are characterized in terms of the existence of a winning strategy in a dialogue game, or better: the non-existence. This can happen when a case is overdetermined (a notion close to logical inconsistency) or underdetermined (there are multiple possible answers with no determined choice). In Gordon's approach to hard cases, there is no dependence on the actual procedure. But in the approach by Hage, Leenes and Lodder, the non-existence of a winning strategy is only one

²⁵ See Miller (2021) for an extensive discussion of the philosophy of justice and John Rawls' analysis of the role of procedure.

of the two parts of the definition of a hard case, as follows. The first part is that the non-existence of a winning strategy shows the necessity of an arational step, i.e., a call to the arbiter. The second part of the definition of a hard case is that such an arational step (in the form of a call to the arbiter) is *in fact* taken by the participants.

Consider now a given decision dialogue that actually has been performed, and that either led to a decision between the parties only, or to a decision that required a call to the arbiter. The approach then leads to four possibilities. First a call to the arbiter was not necessary, and was not made. Then the case was easy. Second a call to the arbiter was not necessary, and was nevertheless made. Then the case was also easy, since an unnecessary call to the arbiter does not make a case hard. Third a call to the arbiter was necessary, and was also made. Then the case was hard. And fourth and finally a call to the arbiter was necessary, but was nevertheless not made. Then the case was easy after all: the participants did not recognize that a call to the arbiter was necessary, and found a solution themselves. Summarizing, only in the third of these four possibilities the case was hard.

Note that the approach by Hage, Leenes and Lodder uses a criterion of hard cases that applies after the decision has been made (after the dialogue has finished), whereas Gordon provides a criterion for hard cases that can be applied before the case goes to court.

7.4 The design of legal decision support systems

The paper contains a discussion of how the procedural perspective on legal decision influences the design requirements for legal decision support systems (Section 3), as follows. A system should allow for decision procedures that can lead to internally consistent, but mutually inconsistent outcomes. Also it should be possible to add or withdraw information during the decision procedure. Such input information might come from outside the system. Furthermore there should be rules of procedure, and, ideally, these should themselves be represented, and thereby possibly become the subject of the discussion. The system should therefore also be able to check whether the procedural rules are followed. Hence, the paper argues against AI systems that autonomously make decisions given certain initial input. Rather such systems should support or guide a discussion between the participants in the procedure. Here the paper also took a position in the discussion in the Netherlands about whether computers could decide legal cases (following the lecture by van den Herik (1991)).

With respect to the design of legal decision support systems for the handling of hard cases, the approach in this paper does not seem to differ much from what Gordon also proposed: both acknowledge that the procedural context should be taken into account, and that non-logical decisions using judicial discretion can be necessary. Also the wider AI and Law discussions at the time recognized the necessity of handling the opposing sides in the setting of a debate. In particular, the HYPO project (Rissland et al., 1984; Rissland and Ashley, 1987; Ashley, 1990) emphasised the process of argumentation involving arguments for both sides, using actual and hypothetical decisions as sources for the arguments and as examples of how to resolve the conflict (see Rissland (2013) for a historical perspective). Another example is the paper by Bench-Capon and Sergot (1988), discussing that

the arguments for and against positions need to be represented when handling open texture.

A particular innovation in the approach by Hage, Leenes and Lodder is that such conflicts between arguments are explicitly modeled in the logical language of DRBL, in a way that aims to do justice to how conflicts are handled in the law. There are sentences expressing that one claim is a reason for or against another, and there are sentences that express whether the reasons for a claim outweigh those against, or the other way around. For instance, in the approach, meta-rules such as *Lex Posterior* and *Lex Superior* may provide reasons for the interpretation of a legal rule, which can be conflicting in a specific example. In such a situation, the approach allows the explicit representation of how the conflict between the reasons is resolved. By its focus on concrete, defeasible, context-dependent logical representations, the paper can be regarded as a precursor of the later research on argumentation schemes (Walton et al., 2008), for which the basis was laid at the 2000 Bonskeid symposium (Reed and Norman, 2004), with applications in law (Bex et al. (2003), Verheij (2003b)²⁶, Atkinson and Bench-Capon (2021)).

In their approach to hard cases, Hage, Leenes and Lodder aim to connect the representations of the logical, the adversarial and the procedural levels of legal argument (cf. the discussion of these levels in Prakken (1995)). Hage later provided a book length presentation of his logical views on the statements about the reasons for and against claims, their weighing and the validity and exclusion of rules and their application (Hage, 1997). Hage (2000) continues the discussion of the role of logic and dialogues in establishing the law. Lodder (1999) expands the approach to a discussion of legal justification using a dialogical approach. Leenes (2001) addresses the handling of burden of proof in a dialogue setting, using the Dutch civil procedure as a case study.

7.5 Final remarks

The paper can be considered as part of the tradition in AI and Law that recognizes that legal reasoning is an interactive process in which decisions are constructed, also and especially when cases are hard (Gardner, 1987; McCarty, 1997; Hafner and Berman, 2002; Bench-Capon and Sartor, 2001). More specifically the paper is one of several contributions to the development of the dialogue perspective in AI and Law in the 1990s, and (considering its citations) has influenced subsequent work on dialogues, argumentation systems and legal reasoning. For instance, the assessment of conflicts in legal reasoning was studied in a dialogue setting by Prakken and Sartor (1996) citing this paper. Vreeswijk (1997) refers to the paper in his study of formal argumentation in which arguments are treated as a kind of defeasible proofs. The Carneades model of argument and burden of proof (Gordon et al., 2007) takes a procedural and dialogical perspective, and connects to Reason-Based Logic (incidentally in a monological version).

The paper was written when Jaap Hage was supervising Arno Lodder, Ronald Leenes and myself during our PhD research, hence the concepts and thoughts in this paper have strongly influenced my research in AI and Law (Verheij, 2020).

²⁶ See section 3 of Araszkievicz et al. (2022) and section 3 of Sartor et al. (2022) elsewhere in this issue for discussions of Bex et al. (2003) and Verheij (2003b) respectively.

In particular, the logical approach to legal reasoning underlying this paper led to the formalization of Reason-Based Logic in my dissertation (abstracting from the dialogue setting), which was applied to the modeling of the relations between rules and principles, and different styles of rule-based analogical reasoning (Verheij et al., 1998). The semi-formal representations of the elements of legal reasoning inspired my work on argumentation schemes (Verheij, 2003b). The perspective on legal decision support systems was one of the inspirations for my work on argument assistance software (Verheij, 2005). Also when Heng Zheng, Davide Grossi and I recently proposed a formal theory of kinds of hard cases (Zheng et al., 2021a), we turned to this paper.

All in all, the paper still stands as a highlight among the 1990s boom of strong papers studying monological and dialogical argumentation in the context of AI and Law. It strikes a rare and excellent balance between careful legal theory, thorough consideration of legal practice, and the explicit modeling of legal decision making. By providing a realistic case study, the paper shows that its approach is not sterile, but applicable to the actual analysis of cases. The paper’s exposition of when a case is hard, and how good lawyers show that a case is hard, provides highly relevant insights for the development of legal technology supporting the hybrid collaboration between humans and machines (Akata et al., 2020).

8 Rationales and argument moves (Loui and Norman, 1995). *Commentary by Bart Verheij*

The paper ‘Rationales and argument moves’ by Ronald Loui and Jeff Norman appeared in the journal in 1995 (Loui and Norman, 1995). The paper has 31 pages and 4 sections. It currently has been cited by 127 scholarly sources.²⁷

8.1 Overview

The paper addresses the topic of rationales in legal decision making, and aims to contribute to the formal explication of various kinds of argument moves that use rationales.

A key idea is that the rationales used in an argumentative dialogue can be interpreted as the summaries (‘compilations’) of extended rationales with more structure. By unpacking such summary rationales, new argument moves are possible, in particular new ways of attacking the argument. For instance, the following dialogue template is in the spirit of the paper:

A: I claim C because of reason R.

B: Unpacking your reasoning, you seem to claim C because of R using additional reason S. I disagree with S, because of reason T. Hence I disagree with your claim C because of R.

A: I agree with your reason T, but I was not using S. Instead I used U, hence my claim C because of reason R.

²⁷ Google Scholar, scholar.google.com, accessed February 19, 2022.

In this format (not appearing in the paper like this), dialogue participant A makes an argument, and participant B unpacks that argument in a certain way, and uses that unpacking in an attack. Subsequently A partially concedes, but disagrees with the unpacking, thereby defending the original position. Below we give more concrete examples.

The paper distinguishes rationales for rules and rationales for decisions. In the authors' terminology, rule rationales express mechanisms for adopting a rule, while decision rationales express mechanisms for forming an opinion about the outcome of a case. As kinds of rule rationale, the authors distinguish compression, specialization and fit (referred to as c-rationales, s-rationales and f-rationales). The kinds of decision rationale are disputation and resolution (d-rationales and r-rationales). Another distinction used is that between object-level and meta-level disputation, where c-, s-, and d-rationales occur in object level disputes, and r- and f-rationales in meta-level discussion. The use of two of these kinds of rationale is illustrated below in sample dialogues. The five rationales are described as examples of 'compilation rationales', that is rationales 'that compile pieces of past disputations, past lines of argument, past preferences of one argument over another, past projections from cases' (p. 160). For each of the five kinds of rationale, structures and attacks are discussed, followed by abstract examples. The paper does not claim or intend a full formalization of rationales and argument moves, but hopes to support further work by 'future automatons of rule-based and case-based legal reasoning' (p. 188).

We illustrate the use of rationales by unpacking a compression rationale and a resolution rationale.

8.2 Unpacking a compression rationale

Here is an example of a small dialogue in which a compression rationale is unpacked, attacked and subsequently defended, following the basic format at the start of this commentary. The unpacking here has the form of adding an intermediate step, thereby interpreting a one step argument as a two step argument. We use a legal example (in the context of Dutch tort law), noting that the original paper focuses on abstract examples involving propositional constants a, b etc. (cf. p. 174 in the paper).

A: I claim that there is no duty to pay the damages ($\neg\text{dut}$) because of the act that resulted in damages (act).

B: Unpacking your reasoning, you seem to claim $\neg\text{dut}$ because of act using the additional intermediate reason that there is a ground of justification (jus). I disagree with jus , because the act was unlawful (unl), so there is no support for jus . Hence there is also no support for $\neg\text{dut}$.

A: I agree with your reason unl and that hence there is no support for jus . But I was not using jus as an intermediate step supporting $\neg\text{dut}$. Instead I used the intermediate step that there was no causal connection between the act and the damages ($\neg\text{cau}$), hence my claim $\neg\text{dut}$ because of act .

A graphical summary of this dialogue is shown in Figure 1. Normal arrows indicate a supporting reason and arrows ending in a cross indicate an attacking reason. All




Fig. 1: Unpacking a compression rationale

abbreviated statements are considered to be successfully supported, except those that are struck-through. Writing the first argument by A as $\text{act} \rightarrow \neg\text{dut}$, B replies in the second move by interpreting the argument as actually having two steps $\text{act} \rightarrow \text{jus} \rightarrow \neg\text{dut}$, and then attacks the unpacked argument on the intermediate step with the argument unl , so that that jus and $\neg\text{dut}$ are no longer supported. But then in the third step A can concede that unl but give the unpacking as $\text{act} \rightarrow \neg\text{cau} \rightarrow \neg\text{dut}$, providing an alternative way to support $\neg\text{dut}$, thereby still maintaining $\text{act} \rightarrow \neg\text{dut}$.

8.3 Unpacking a resolution rationale

Again we give a mini-dialogue, illustrating how the authors approach the idea of resolution rationales and argument attack. The example unpacks an argument as using the weighing of two conflicting reasons.

A: I claim that there is a duty to pay the damages (dut) because of the act that resulted in damages (act).

B: Unpacking your reasoning, you seem to claim dut because of act using the weighing of two reasons, one for the duty to pay (the high probability of damages, prb) and one against (the mild nature and low scale of the possible damages, ntr). I disagree with this weighing, because the nature and scale of the possible damages here was exceptionally low and so outweighs the high probability of damages. Hence I disagree with your claim dut because of act .

A: I agree with your weighing of the two reasons you mention. But I was using an additional reason for the duty to pay (it was easy to take precautionary measures, mea), and the two reasons for the duty to pay taken together ($\text{prb} \wedge \text{mea}$) outweigh the one reason against (ntr). Hence my claim dut because of act .

A graphical summary is shown in Figure 2. If we write A's argument in the first dialogue move as $\text{act} \rightarrow \text{dut}$, then in the second move B unpacks A's reasoning by claiming that there are reasons for and against dut , namely prb and ntr , but whereas A had preferred prb , B claims that $\text{prb} < \text{ntr}$. According to B's weighing of these two reasons, the conclusion should therefore be $\neg\text{dut}$. In the third step, A agrees with B's weighing of the two reasons, but adds a third reason ($\text{mea} \rightarrow \text{dut}$) that turns the outcome to the other side, since $\text{prb} \wedge \text{mea} > \text{ntr}$, reinstating dut .




Fig. 2: Unpacking a resolution rationale

8.4 Later developments

This paper was published in the middle of the 1990s, in a period of especially active research efforts investigating the formal structure of argumentation with rules and cases, by taking inspiration from the law. The paper refers for instance to work available then by Ashley, Branting, Gordon, Hage, Pollock, Prakken, Rissland, Sartor, Simari and Vreeswijk (see for instance Ashley (1990), Branting (1993), Gordon (1995), Hage et al. (1993), Pollock (1992), Prakken and Sartor (1996), Rissland (1983), Simari and Loui (1992) and Vreeswijk (1997)).

Since the paper appeared, research on computational argumentation and the modeling of reasoning with rules and cases in law has continued (for references and overviews see e.g. Atkinson et al. (2017), Baroni et al. (2020), Prakken and Sartor (2015) and Verheij (2020)). We briefly mention three relevant developments: argument structure, argumentation schemes and argumentation semantics.

First, argument structure has been widely studied since the paper appeared. In particular, the argument structure and its evaluation as suggested in Figure 1 can be reconstructed in various proposals for structured arguments and their evaluation, such as DefLog (Verheij, 2003a). An alternative approach to representing argument structure is ASPIC+ (Prakken, 2010).

Second, there is connection between unpacking and argumentation schemes (Walton (1996) and Walton et al. (2008)). Argumentation schemes can be thought of as concrete, defeasible reasoning patterns that can be used to suggest moves in an argumentative dialogue. Representation formats for argumentation schemes (e.g., Verheij (2003b) and Atkinson and Bench-Capon (2007)) include critical questions which can be used to object to the argument. Thus, in the compression rationale *Isn't the act unlawful?* could be seen as a critical question to A's argument. Note however, that this scheme does not uncover the structural relation that B considers to be the problem of A's argument, namely that A has used a problematic intermediate step via *jus*. One could say that the scheme allows B to add a possible counterargument (*un1*), but not to unpack the chain from *act* to *-dut*.

Third we mention the connection between argumentation semantics and unpacking. The paper was published in the same year as Dung's influential work on the evaluation of sets of argument attack relations (Dung, 1995), that soon became the standard reference for all work on the semantics of argumentation (see Baroni et al. (2020) for a historical perspective at the occasion of Dung's paper's 25th anniversary). In Dung's work arguments are entirely abstract and so the nature

of the attacks is hidden. But the use of rationales shows that the arguments may comprise a number of steps, and which step is attacked may be important. DefLog (Verheij, 2003a) and ASPIC+ (Prakken, 2010) provide different means of locating attacks in terms of the argument structure. The formalism in Verheij (2017) uses a different approach to the semantics of arguments, by defining notions of argument validity in such a way that both a packed and related unpacked arguments are valid.

8.5 Final remarks

The paper showed at the time of publication, and still shows now, a playful and creative use of formal syntax, focusing on ideas, and not so much on mathematical constraints and definitions. Also its scholarly writing style remains distinctive in its independence. The high level general message of the paper stands: arguments can be combined, summarized, packed and then again unpacked in argumentative dialogue, and that formal syntax can help for analyzing such compilation and decompilation processes. Considering later research about these and similar ideas, there seems to be little work that directly connects to compilation and decompilation. References to the paper typically mention the process of unpacking, or to its cataloguing of various structures of argument moves. We suggested that later work on argument structure, argumentation schemes and argumentation semantics is relevant for the investigation of unpacking arguments. Still it seems that the formal, syntactic connections between arguments, rules and cases are not yet ready to fully address argument compilation and decompilation in its various forms as discussed by Loui and Norman, and that further work would be valuable.

9 Modelling reasoning with precedents in a formal dialogue game (Prakken and Sartor, 1998). *Commentary by Trevor Bench-Capon*

Although Skalak and Rissland (1992) (see Section 2) showed that knowledge of the rules governing the domain in question and knowledge of cases to enable interpretation of these rules were both essential to a complete account of legal reasoning, the style of reasoning with rules remained different from the style of reasoning with cases. The two were intertwined, but not unified. One major contribution of Prakken and Sartor (1998) was that it provided a way of representing precedents as rules, allowing a uniform representation of both case and statute knowledge. There is much else in the paper, but this means of representing precedents as rules remains the standard way of incorporating precedents into logic based approaches (e.g. Bench-Capon and Atkinson (2021)), and was a key influence on the series of current formal accounts of precedential reasoning initiated by Horty (Horty, 2011)²⁸, summarised in Prakken (2021). Before describing the representation of precedents as rules in detail, however, we should provide some context to the paper.

Throughout the 90s a widespread method of representing the adversarial and procedural aspects of legal reasoning was through the specification of dialogue protocols and games. 1993 saw the publication of both Gordon (1993) (see Section 6)

²⁸ See Section 10 of Sartor et al. (2022), in this issue.

and Hage et al. (1993) (see Section 7). The former used its dialogue to identify contested issues, but more typical of what was to come was Hage et al. (1993), which modelled decision making in hard cases. The various dialogue games had a number of different foundations. Some were based on particular logics: Gordon (1993) used Conditional Entailment (Geffner and Pearl, 1992) while Hage et al. (1993) used a dialectical version of Reason Based Logic (Hage, 1993). In Bench-Capon et al. (2000) two dialogue games were presented: one based on the existing dialogue game DC (Mackenzie, 1979) and another based on the argument scheme of Toulmin (Toulmin, 1958). Prakken and Sartor had proposed a dialogue game for resolving normative conflicts in Prakken and Sartor (1996) based on the argumentation semantics for logic programming of Dung (1993)²⁹. Prakken and Sartor (1998) uses this dialogue game to model precedents.

The dialogue game presented in Prakken and Sartor (1996) differs from the other games mentioned above in that it is intended only to provide a proof theory for an argumentation-based semantics. This contrasts with the intention to model legal procedures (Gordon, 1993) or to regulate debates between human and/or artificial agents as in Hage et al. (1993) and Bench-Capon et al. (2000). In Prakken and Sartor (1996) each side takes it in turn to put forward an argument. After an initial argument has been presented, the opponent presents an argument attacking that argument (as in the two party immediate response games found in the argumentation literature (Vreeswijk and Prakken (2000) and Dunne and Bench-Capon (2003)). The game terminates when no moves are possible, the winner being the player to make the last move. An argument is acceptable if the proponent has a winning strategy. In Dung's framework and Dunne and Bench-Capon (2003) all attacks succeed as defeats, but Prakken and Sartor (1996) allows priorities between arguments, so that a stronger argument may resist the attack of a weaker argument. The dialogue game in Prakken and Sartor (1996) is first presented with fixed priorities and then extended to allow arguments about priorities, so that legal principles such as *Lex Superior*, *Lex Specialis* and *Lex Posterior* can be represented, and priorities between rules derived within the system. In Prakken and Sartor (1998), with its focus on case-based reasoning, these priorities come from decisions in precedent cases rather than legal principles. Prakken and Sartor (1998) also extends the original game by modelling some case based reasoning moves (broadening and distinguishing) to enable the participants to introduce new premises. In this short section, it is impossible to do justice to all aspects of the paper, and so I will focus specifically on an account of how precedents can be represented.

An argument is a sequence of ground instances of *rules*, where a rule comprises a head representing the conclusion, and a body which is a conjunction of rule consequents representing the antecedent. Facts are represented by rules without antecedents. Given that arguments rely on rules, if precedents are to be used as the basis of arguments, they need to be represented as rules of this form. In CATO (Aleven and Ashley, 1995), cases are associated with a set of factors, some pro-plaintiff and some pro-defendant, and an outcome. The pro-plaintiff factors offer reasons to find for the plaintiff and the pro-defendant factors offer reasons to find

²⁹ This paper was a precursor of Dung's abstract argumentation frameworks (Dung, 1995). Prakken subsequently developed a dialogue game based on those frameworks in Prakken (1997b).

for the defendant. Now, if we have a decided case, C , containing pro-precedent factors P and pro-defendant factors D , then the conjunction of all factors in P will be the strongest³⁰ reason to decide C for the plaintiff and the conjunction of all factors in D the strongest reason to decide C for the defendant. The outcome in the case will show which of these two reasons is stronger. This means we have three rules:

- r1 $P \rightarrow \textit{plaintiff}$;
- r2 $D \rightarrow \textit{defendant}$;
- r3 $C \rightarrow r2 \prec r1$ if the decision was for the plaintiff and $C \rightarrow r1 \prec r2$ if the decision was for the defendant.

This representation sees precedents as providing a one step argument from factors to outcome, which was the view taken in subsequent approaches such as Bench-Capon (1999) and the formalisations of precedential constraint stemming from Horty (2011). Prakken and Sartor (1998), however, argue strongly that precedents should be seen in terms of multi-step arguments. Often the importance of precedent will be with respect to a particular issue in the case (Branting, 1991). Thus if we partition the factors according to the issues of the case, using, for example, the abstract factor hierarchy of Alevén (2003), we can get a finer grained representation of the argument. We now represent the case as $P_1 \cup D_1 \cup \dots \cup P_n \cup D_n$, where P_i are the pro-plaintiff factors relating to issue i and D_i are the pro-defendant factors relation to issue i . We can now produce a set of three rules for each *issue*:

- r4 $P_i \rightarrow I_i^P$, where I_i^P means that issue i is found for the plaintiff;
- r5 $D_i \rightarrow I_i^D$; where I_i^D means that issue i is found for the defendant;
- r6 $C \rightarrow r5 \prec r4$ if the issue was found for the plaintiff in C and $C \rightarrow r4 \prec r5$ if the issue was found for the defendant in case C .

We now write a set of three rules, using the issues in the antecedents to show how the issues determine the outcome. Suppose we have a case with three issues, of which two were found for the plaintiff and one for the defendant but the defendant won the case. This would give the rules:

- r7 $I_1^P \wedge I_2^P \rightarrow \textit{Plaintiff}$
- r8 $I_3^D \rightarrow \textit{Defendant}$
- r9 $C \rightarrow r7 \prec r8$

Not only does this more faithfully reflect the reasoning in the case, but it has the practical advantage that an inference is not blocked by a distinction which is irrelevant because it pertains to a different issue. This two step reasoning, from factors to issues and then from issues to outcome was later used in IBP (Brüninghaus and Ashley, 2003) and Grabmair’s VJAP (Grabmair, 2017). More recently it has been argued that adopting this finer grained representation would improve the formal accounts of precedential constraint (Bench-Capon and Atkinson, 2021). Even

³⁰ This assumes that the conjunction of two factors favouring the same side will always be stronger than the factors individually. This assumption is queried in Prakken (2005), where an apparent counter example is given. However, such situations can be avoided by modelling the domain differently, using different factors for which the original “factors” are facts (e.g. Horty and Bench-Capon (2012), footnote 17). Arguably it is a necessary feature of factors that they always favour a particular side, and this should hold whatever the context set by other factors (Bench-Capon, 2017).

finer granularity would be possible, to give rise to three step arguments, but that will often associate too few factors with each sub-issue to be useful.

If the new case has an exact match with plaintiff and defendant rules for which a preference exists, it can be decided from the precedents. The same is true when any differences, such as when the new cases has an additional factor for the winning side, only strengthen the case for that outcome. Often, however, the available precedent cases contain an additional factor for the winning side, so that the rule does not match, or the new case contains an additional factor for the losing side, so that there is a stronger reason for the other side, for which no preference can be shown. This means that there will be no way, based on existing precedents, to choose between them. This gives a safe model of precedential constraint in terms of factors, but does not constrain as many cases as one might wish. Using a finer grained representation improves the situation, but the problem remains.

To allow reasoning about cases not constrained by the precedents Prakken and Sartor (1998) uses two argument moves taken from cased based reasoning approaches such as CATO and CABARET (See Section 2), to provide heuristics for introducing new premises that case extend the current theory. These moves are *rule broadening* and *distinguishing*.

Rule broadening can be used when no rule matches but some rule fails on only one of the literals in the antecedent. Although this does not seem a reasonable move in the case of rules derived from statutes, it is a sensible move for the rules derived from precedents, since not all the reasons may have been needed to defeat the opponent. Consider a precedent with factors $P \cup D$, found for the plaintiff. Now although we can be certain that P is a stronger reason than D , it may well have been that a *subset* of P would also have been sufficient to defeat D . If this is so, we can prefer the rule $P_1 \rightarrow \textit{plaintiff}$ where P_1 is a subset of P which does hold for the current case.

Distinguishing is where we have a precedent $P \rightarrow \textit{plaintiff}$ as above, but this time the current case contains an additional pro-defendant factor d . Now although P is a stronger reason than D , the reason in the precedent, it is arguably not as strong as the reason available in the new case $D \cup d$. Thus introducing this additional rule for the defendant will enable an attack on the plaintiff rule not available without so extending the theory.

The additional premises introduced by rule broadening and distinguishing are termed *introducibles* in Prakken and Sartor (1998). They do note, however, that the resolution of a dispute involving introducibles requires a decision from an arbiter or judge to determine whether the broadened rule does in fact remain stronger than the opposing rule, or the distinction is significant³¹. Whereas the *a fortiori* arguments show the new case to be constrained, arguments with introducibles may be accepted, and so become precedents for future use, but they may also be rejected, leaving the case law unchanged.

This method of modelling precedents provided the inspiration for Bench-Capon (1999), with plaintiff and defendant reasons represented as two separate partial orders and precedents as ordering links between them. If a new case cannot link the orders in a way which does not introduce a cycle, it is constrained by the precedent base. These notions were made rigorous in Horty (2011) and Horty

³¹ Such an appeal to an arbiter is a defining notion of a “hard case” in Hage et al. (1993), discussed in Section 7.

and Bench-Capon (2012)³². Here the basic notions of Prakken and Sartor (1998), without rule broadening, are termed the *results* model, which Horty attributes to Larry Alexander (Alexander, 1989). To replace rule broadening, Horty advocates the use of the *reason* model, which he attributes to Greg Lamond (Lamond, 2005). In the reason model the precedent is itself held to give a rule with a subset of the factors for the winning side. How this subset is determined is not stipulated by Horty, but the suggestion is that this can be discovered from reading the decision, perhaps by adopting the *ratio decendi* of the case. Distinguishing is then the way of showing that case is not constrained by the precedent base. Horty’s account is not seen as a dialogue, but as a question of whether or not the decision in a new case is constrained, and so no additional premises can be introduced, which is why rule broadening and distinguishing are built into his representation and use of the rules. For a formal comparison of the results and the reason models see Prakken (2021).

In Prakken and Sartor (1998) no argument is given for whether or not an inductible should be accepted by the arbiter. In Bench-Capon and Sartor (2003) the suggestion of Berman and Hafner (1993)³³ was applied to the rule based representation of precedents. Here the preferences revealed by precedents were held to apply not to the sets of *factors*, but to the *values* promoted or demoted by following the pro-plaintiff or pro-defendant rules. Because several different sets of factors ground a rule promoting the same value, this preference is applicable to many more cases than those using the precise reasons of the precedents.

The dialectical principles of Prakken and Sartor (1998) were further developed by Prakken into the abstract framework for representing structured arguments, ASPIC+ (Prakken (2010) and Modgil and Prakken (2014)). Here the tree of arguments and counter arguments generated by the dialogues of Prakken and Sartor (1998) become the subarguments and attacks on them found in ASPIC+. ASPIC+ has been widely used in general AI argumentation (e.g. Modgil and Prakken (2013)) and applied to AI and Law in Prakken et al. (2015) and Prakken (2019).

There are several highly influential ideas in Prakken and Sartor (1998). For the general field of AI argumentation it can be seen as the forerunner of ASPIC+ (Prakken, 2010), which provides a framework for structured argumentation intended to generate arguments that can then be evaluated using Dung’s abstract framework, providing a link between the argument structure essential for the modelling of real disputes, and the formal properties of abstract argumentation. For AI and Law in particular its method of representing precedents as rules can be seen as providing the key means to unify rule and case based reasoning, and the foundations for current formal accounts of precedential reasoning.

10 Hybrid rule - neural approach for the automation of legal reasoning in the discretionary domain of family law in Australia (Stranieri et al., 1999). Commentary by Matthias Grabmair.

The 1999 paper by Stranieri, Zeleznikow, Gawler, and Lewis contributes a case study on a system called “Split-Up” that models discretionary property division

³² See Section 10 of Sartor et al. (2022), elsewhere in this issue.

³³ See Section 2 of Sartor et al. (2022), elsewhere in this issue.

decisions in divorce proceedings in Australian Family law by means of argument chains, rule systems, and neural network implementations. The domain is formalized into a set of 94 factors that form the basis of a consultation dialogue with the user guided by an argument chain traversing all relevant case information at the end of which the system predicted an asset split percentage. Factors are organized in a hierarchy and allow the derivation of higher-level legal information up to the overall determination of the outcome. The authors model each aggregative inference by either using rules (including transition networks) or by crafting a neural network. Their knowledge representation methodology in making the choice between the two for each higher-level concept centers around classifying it along two dimensions:

- open-textured-ness versus well-defined-ness (i.e., the degree to which a legal term encompasses a potentially large set of possible fact patterns); and
- bounded-ness versus unbounded-ness (i.e., the extent to which all relevant legal concerns for a decision about the concept can be exhaustively enumerated by an expert).

All neural networks used a single hidden layer and provide a minimal machine learning inference mechanism between the values of the legal concepts forming input and output in the respective case. Their parameters were trained using real historical family law data and the system’s split percentage outputs are compared with those produced by experts in a small evaluation experiment. Split Up’s contribution to the field lies in its nuanced composition of expert-crafted rules and neural networks trained from data that is faithful to a factor-based model of Australian family law, and which forms an early example of transparent and explainable use of machine learning techniques in legal decision support systems. It is cited as among the pioneer applications of neural networks in modelling legal inference (e.g., Steging et al. (2021)). Xu and Yuan (2009) emphasize Split-Up’s extension of purely rule-based modeling and explain it in the same context as hybrid systems extending rules with case-based reasoning Skalak and Rissland (1992)³⁴.

Split Up was a long running project and the 1999 paper is the culmination of a series of publications that began with Zeleznikow and Stranieri (1995) and Zeleznikow et al. (1995) and continued with Zeleznikow and Stranieri (1997), Zeleznikow and Stranieri (1998), and Stranieri and Zeleznikow (1998). It influenced several other projects in the field of decision and negotiation support. For example, Mackenzie et al. (2015) developed an experimental system to support criminal law sentencing and guilty plea negotiations, explicitly taking inspiration from Split Up and employing neural networks as part of their system. Miah et al. (2020) cite it as an example of explainable decision support in their work on a prototype clinical decision support system focusing on sleep disorder diagnosis and treatment. In follow up work, Yearwood and Stranieri (2006) formalize the notion of a “generic/actual argument model” (GAAM), which they describe as underlying Split Up. The formalism derives from the Toulmin argument model (Toulmin, 1958) and focuses on parties instantiating “actual” arguments from “generic templates” by means of a variable-slot-value representation. Notably, the GAAM does not force the use of particular inference procedures, thereby allowing Split Up’s concept-by-concept choice between rules or neural networks. Moulin et al. (2002)

³⁴ See the discussion in Section 2.

recognized this as having been “left implicit in Toulmin’s formulation: an inference procedure, algorithm or method used to infer an assertion from datum”.

Given the prominence of deep learning in the current landscape of AI and Law research, Split Up naturally stands out as among the first uses of neural networks in the field. Even earlier, Bochereau et al. (1991) trained a neural network on a small collection of French Council of State opinions and extracted decision rules from the trained model via a combinatorially generated validation dataset. Bench-Capon (1993) explored the use of basic hidden layer models on synthetic data to predict outcomes and derive input feature weights, albeit with mixed results. Notably, these works focused on the application of neural networks to identify how influential certain information is for the decision (i.e., rules and feature importances) and did not engage in comparative benchmarking towards the best prediction performance, as is common today. Their input variables directly corresponded to decision-relevant information elements of legal case configurations, and the experiments centered around exploring whether and how neural networks could model legal reasoning. This marks another difference to modern deep learning, where models are intended to emulate human cognition on as raw as possible data with very little expert-informed feature engineering. A critical survey published during the same period (Hunter, 1999) argued that legal reasoning is not akin to sensory pattern recognition, and that the use of neural networks for models of legal decision making is to be scrutinized against relevant legal theory. It points out the lack of ability of statistical methods to justify their predictions in the same satisfactory way as it is possible in symbolic inference. Additionally, the limited availability of suitable, sufficiently large, datasets constrained the scope of neural network research for modeling legal inference at that time.

These obstacles may be among the reasons why the use of neural networks in published AI and Law work declined sharply between 1999 and 2017. To some degree, this appears to correlate with the general trend in machine learning in that period, which saw widespread use of other model architectures, like tree- and forest-based models as well as kernel-driven support vector machines. Neural networks continued to be used for cognitive tasks, but it was not until around 2012, when large convolutional networks trained on graphics processors achieved breakthrough performance in computer vision, that neural models quickly became the dominant technology. Shortly after, recurrent neural architectures, like the long-short-term-memory network (or LSTM (Hochreiter and Schmidhuber, 1997)), and soon attention-based models (e.g. Du and Huang (2018)), similarly took over general natural language processing. This development eventually reached the AI and Law community and led to a surge in the use of such models for the analysis of legal text from around 2017 onwards. Early works include, for example, contract element extraction (Chalkidis et al., 2017), sentence modality classification (Neill et al., 2017), as well as requisite and effectuation segment recognition in statutory text (Nguyen et al., 2018)³⁵.

Notably, this revival of neural models in AI and Law did not resume modeling legal reasoning as tackled by the models in the 1990s, but has rather been focusing on tasks where they are most suitable: cognition-like processing of textual data with minimal feature engineering on large corpora, enabled by the increasing number of publicly available textual datasets for legal NLP research. As in mainstream

³⁵ See section 5 of Villata et al. (2022), elsewhere in this issue.

NLP, transfer learning from large language models has become the dominant approach (see, e.g., Zheng et al. (2021b)). Overall, one can observe increased focus on benchmarking models along quantitative metrics, typically with regard to high level classification/prediction goals such as case outcome variables and document-level keywords (e.g., Chalkidis et al. (2022)). This forms a contrast to Split Up’s use of NNs as specialized components that model discretionary decisions via statistically trained parameters in an otherwise symbolic decision process.

Some recent work re-engages with the idea of combining NNs with structural models of legal inference and argumentation. For example, Holzenberger and Van Durme (2021) apply deep neural NLP models to identify argument slots from the text of legal rules (e.g., the “taxpayer”, “income”, etc.) and find suitable filling elements from fact descriptions (e.g., names of persons, monetary amounts). The instantiated schemas allow the derivation of an answer to a question about a case that is coupled with the elements of the legal provision, and the black-box inference of the NN model is predominantly used for cognition-like functions, such as identifying concept-mentions in the statute and unifying them with entity-mentions in the case text. This facilitates explainability and aligns with the traditional paradigm of modeling legal reasoning as argumentative inference in synergy with statistical learning.

In conclusion, Stranieri et al. (1999) and the Split Up system form an important example of an early real-life application using a Toulmin-inspired argument model. While its influence as an early adopter of neural networks is limited, its implementation of judicial discretion by means of small network components remains a reference point along the path to more advanced explainable AI and Law applications.

References

- Zeynep Akata, Dan Balliet, Maarten De Rijke, Frank Dignum, Virginia Dignum, Guszti Eiben, Antske Fokkens, Davide Grossi, Koen Hindriks, Holger Hoos, et al. A research agenda for Hybrid Intelligence: Augmenting human intellect with collaborative, adaptive, responsible, and explainable Artificial Intelligence. *Computer*, 53(08):18–28, 2020.
- Latifa Al-Abdulkarim, Katie Atkinson, and Trevor Bench-Capon. A methodology for designing systems to reason with legal cases using ADFs. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
- Carlos E Alchourrón. Conditionality and the representation of legal norms. In Antonio A. Martino and Fiorenza Socci-Natali, editors, *Automated Analysis of Legal Texts*, pages 175–186. North Holland, 1986.
- Carlos E Alchourrón and David Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.
- Carlos E Alchourrón and David Makinson. On the logic of theory change: Safe contractions. *Studia Logica*, 44:405–422, 1985.
- Carlos E Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

- Vincent Aleven. Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment. *Artificial Intelligence*, 150(1-2):183–237, 2003.
- Vincent Aleven and Kevin D Ashley. An instructional environment for practicing argumentation skills. In *Proceedings of AAAI 1994*, pages 485–492, 1994.
- Vincent Aleven and Kevin D Ashley. Doing things with factors. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 31–41, 1995.
- Larry Alexander. Constrained by precedent. *Southern California Law Review*, 63:1, 1989.
- Robert Alexy. *A Theory of Legal Argumentation*. Clarendon Press, (1st edition in German 1978), 1989.
- Michał Araszkiwicz, Trevor Bench-Capon, Enrico Francesconi, Marc Lauritsen, and Antonino Rotolo. Thirty years of AI and Law: Overviews. *Artificial Intelligence and Law*, 30(4):This issue, 2022.
- Julieth Patricia Castellanos Ardila, Barbara Gallina, and Guido Governatori. Compliance-aware engineering process plans: the case of space software engineering processes. *Artificial Intelligence and Law*, 29(4):587–627, 2021.
- Kevin D Ashley. *Modeling legal arguments: Reasoning with cases and hypotheticals*. MIT press, 1990.
- Kevin D Ashley and Vincent Aleven. A computational approach to explaining case-based concepts of relevance in a tutorial context. In *Proceedings of the 4th Workshop on Case-Based Reasoning*, pages 257–268, 1991.
- Kevin D Ashley and Stefanie Brüninghaus. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*, 17(2):125–165, 2009.
- Katie Atkinson, editor. *Artificial Intelligence and Law: Special Issue on Modelling Popov v Hayashi*, volume 20:1. 2012.
- Katie Atkinson and Trevor Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 171:855–874, 2007.
- Katie Atkinson and Trevor Bench-Capon. Argumentation schemes in AI and Law. *Argument and Computation*, 12(3):417–434, 2021.
- Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Simari, Matthias Thimm, and Serena Villata. Toward artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.
- Pietro Baroni, Francesca Toni, and Bart Verheij. Introduction to the special issue ‘On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games: 25 years later’. *Argument and Computation*, 11(1–2):1–14, 2020.
- Trevor Bench-Capon. Deep models, normative reasoning and legal expert systems. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, pages 37–45, 1989.
- Trevor Bench-Capon. Neural networks and open texture. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 292–297, 1993.
- Trevor Bench-Capon. Some observations on modelling case based reasoning with formal argument models. In *Proceedings of the 7th International Conference on Artificial Intelligence and Law*, pages 36–42, 1999.
- Trevor Bench-Capon. Representing Popov v Hayashi with dimensions and factors. *Artificial Intelligence and Law*, 20(1):15–35, 2012.

- Trevor Bench-Capon. HYPO's legacy: introduction to the virtual special issue. *Artificial Intelligence and Law*, 25(2):205–250, 2017.
- Trevor Bench-Capon and Katie Atkinson. Precedential constraint: The role of issues. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pages 12–21, 2021.
- Trevor Bench-Capon and Frans Coenen. Isomorphism and legal knowledge based systems. *Artificial Intelligence and Law*, 1(1):65–86, 1992.
- Trevor Bench-Capon and Paul Dunne. Argumentation in Artificial Intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.
- Trevor Bench-Capon and Justin Forder. Knowledge representation for legal applications. In Trevor Bench-Capon, editor, *Knowledge-Based Systems and Legal Applications*, pages 245–263. Academic Press, 1991.
- Trevor Bench-Capon and Thomas F Gordon. Isomorphism and argumentation. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, pages 11–20, 2009.
- Trevor Bench-Capon and Sanjay Modgil. Norms and value based reasoning: justifying compliance and violation. *Artificial Intelligence and Law*, 25(1):29–64, 2017.
- Trevor Bench-Capon and Henry Prakken. Using argument schemes for hypothetical reasoning in law. *Artificial Intelligence and Law*, 18(2):153–174, 2010.
- Trevor Bench-Capon and Giovanni Sartor. Theory based explanation of case law domains. In *Proceedings of the 8th International Conference on Artificial intelligence and Law*, pages 12–21, 2001.
- Trevor Bench-Capon and Giovanni Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1-2):97–143, 2003.
- Trevor Bench-Capon and Marek Sergot. Towards a rule-based representation of open texture in law. In C. Walter, editor, *Computer Power and Legal Language: The Use of Computational Linguistics, Artificial Intelligence, and Expert Systems in the Law*, pages 39–60. Quorum Books, New York, 1988.
- Trevor Bench-Capon, Gwen Robinson, Tom Routen, and Marek Sergot. Logic programming for large scale applications in law: A formalisation of Supplementary Benefit legislation. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 190–198, 1987.
- Trevor Bench-Capon, Tim Geldard, and Paul Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8(2):233–254, 2000.
- Donald H Berman. Developer's choice in the legal domain: the Sisyphean journey with CBR or down hill with rules. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 307–309, 1991.
- Donald H. Berman and Carole L. Hafner. Indeterminacy: A challenge to logic-based models of legal reasoning. In *Yearbook of Law Computers and Technology Volume 3*, pages 1–35. 1987.
- Donald H Berman and Carole L Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 50–59, 1993.
- Floris Bex, Henry Prakken, Chris Reed, and Douglas Walton. Towards a formal account of reasoning about evidence: argumentation schemes and generalisations. *Artificial Intelligence and Law*, 11(2):125–165, 2003.

- Hanif Bhuiyan, Guido Governatori, Andy Bond, Sebastien Demmel, Mohammad Badiul Islam, and Andry Rakotonirainy. Traffic rules encoding using De-feasible Deontic Logic. In *Proceedings of JURIX 2020*, pages 3–12, 2020.
- Carlo Biagioli, Paola Mariani, and Daniela Tiscornia. ESPLEX: A rule and conceptual model for representing statutes. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 240–251, 1987.
- Laurent Bochereau, Danièle Bourcier, and Paul Bourguine. Extracting legal knowledge by means of a multilayer neural network application to municipal jurisprudence. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 288–296, 1991.
- L Karl Branting. Reasoning with portions of precedents. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 145–154, 1991.
- L Karl Branting. A computational model of ratio decidendi. *Artificial Intelligence and Law*, 2(1):1–31, 1993.
- L Karl Branting, Craig Pfeifer, Bradford Brown, Lisa Ferro, John Aberdeen, Brandy Weiss, Mark Pfaff, and Bill Liao. Scalable and explainable legal prediction. *Artificial Intelligence and Law*, 29(2):213–238, 2021.
- Gerhard Brewka. *Nonmonotonic Reasoning. Logical Foundations of Commonsense*. Cambridge University Press, Cambridge, 1991.
- Stefanie Brüninghaus and Kevin D Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 233–242. ACM, 2003.
- Brigitte Burgemeestre, Joris Hulstijn, and Yao-Hua Tan. Value-based argumentation for justifying compliance. *Artificial Intelligence and Law*, 19(2):149–186, 2011.
- Ilias Chalkidis, Ion Androutsopoulos, and Achilleas Michos. Extracting contract elements. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law*, pages 19–28, 2017.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. LexGLUE: A benchmark dataset for legal language understanding in English. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, 2022.
- Frans Coenen and Trevor Bench-Capon. Exploiting isomorphism: development of a kbs to support british coal insurance claims. In *Proceedings of the 3rd International conference on Artificial Intelligence and Law*, pages 62–68, 1991.
- Frans Coenen and Trevor Bench-Capon. Building knowledge based systems for maintainability. In *Proceedings of DEXA 1992*, pages 415–420. Springer, 1992.
- Paolo Di Giusto and Guido Governatori. A new approach to base revision. In Pedro Barahona and José Júlio Alferes, editors, *Progress in Artificial Intelligence, 9th Portuguese Conference on Artificial Intelligence*, volume 1695 of *Lecture Notes in Computer Science*, pages 327–341. Springer, 1999.
- Changshun Du and Lei Huang. Text classification research with attention-based recurrent neural networks. *International Journal of Computers Communications and Control*, 13(1):50–61, 2018.
- Phan Minh Dung. An argumentation semantics for logic programming with explicit negation. In *Proceedings of the 10th International Conference on Logic Programming*, pages 616–630, 1993.

- Phan Minh 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.
- Paul Dunne and Trevor Bench-Capon. Two party immediate response disputes: Properties and efficiency. *Artificial Intelligence*, 149(2):221–250, 2003.
- Anne von der Lieth Gardner. *An Artificial Intelligence approach to legal reasoning*. MIT press, 1987.
- Hector Geffner and Judea Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2-3):209–244, 1992.
- Seth R Goldman, Michael C Dyer, and Margot Flowers. Precedent-based legal reasoning and knowledge acquisition in contract law: A process model. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 210–221, 1987.
- Thomas F Gordon. An abductive theory of legal issues. *International Journal of Man-Machine Studies*, 35:95–118, 1991.
- Thomas F Gordon. The Pleadings game. *Artificial Intelligence and Law*, 2(4): 239–292, 1993.
- Thomas F Gordon. *The Pleadings game: An Artificial Intelligence Model of Procedural Justice*. Springer, 1995.
- Thomas F Gordon. Introducing the Carneades Web application. In *Proceedings of the 14th International Conference on Artificial Intelligence and Law*, pages 243–244, 2013.
- Thomas F Gordon and Douglas Walton. A Carneades reconstruction of Popov v Hayashi. *Artificial Intelligence and Law*, 20(1):37–56, 2012.
- Thomas F Gordon and Douglas Walton. Formalizing balancing arguments. In *Proceedings of COMMA 2016*, pages 327–338, 2016.
- Thomas F Gordon, Henry Prakken, and Douglas Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10-15):875–896, 2007.
- Guido Governatori. Representing business contracts in RuleML. *International Journal of Cooperative Information Systems*, 14(2-3):181–216, 2005.
- Guido Governatori. Thou shalt is not you will. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 63–68, 2015.
- Guido Governatori and Francesco Olivieri. Unravel legal references in defeasible deontic logic. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pages 69–78, 2021.
- Guido Governatori, Vineet Padmanabhan, Antonino Rotolo, and Abdul Sattar. A defeasible logic for modelling policy-based intentions and motivational attitudes. *Logic Journal of the IGPL*, 17(3):227–265, 2009.
- Guido Governatori, Florian Idelberger, Zoran Milosevic, Régis Riveret, Giovanni Sartor, and Xiwei Xu. On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artificial Intelligence and Law*, 26(4):377–409, 2018.
- Matthias Grabmair. *Modeling Purposive Legal Argumentation and Case Outcome Prediction using Argument Schemes in the Value Judgment Formalism*. PhD thesis, University of Pittsburgh, 2016.
- Matthias Grabmair. Predicting Trade Secret case outcomes using argument schemes and learned quantitative value effect tradeoffs. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law*, pages 89–98, 2017.

- Carole D Hafner and Donald H Berman. The role of context in case-based legal reasoning: teleological, temporal, and procedural. *Artificial Intelligence and Law*, 10(1):19–64, 2002.
- Jaap Hage. Monological Reason-Based Logic: A low level integration of rule-based reasoning and case-based reasoning. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 30–39, 1993.
- Jaap Hage. A theory of legal reasoning and a logic to match. *Artificial Intelligence and Law*, 4(3):199–273, 1996.
- Jaap Hage. *Reasoning with Rules. An Essay on Legal Reasoning and Its Underlying Logic*. Kluwer Academic Publishers, 1997.
- Jaap Hage. Dialectical models in Artificial Intelligence and Law. *Artificial intelligence and Law*, 8(2):137–172, 2000.
- Jaap Hage, Ronald Leenes, and Arno R Lodder. Hard cases: a procedural approach. *Artificial Intelligence and Law*, 2(2):113–167, 1993.
- Mustafa Hashmi and Guido Governatori. Norms modeling constructs of business process compliance management frameworks: a conceptual evaluation. *Artificial Intelligence and Law*, 26(3):251–305, 2018.
- Henning Herrestad. Norms and formalization. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 175–184, 1991.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Nils Holzenberger and Benjamin Van Durme. Factoring statutory reasoning as language understanding challenges. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2742–2758, 2021.
- John F Horty. Reasons and precedent. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 41–50, 2011.
- John F Horty and Trevor Bench-Capon. A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20(2):181–214, 2012.
- Dan Hunter. Out of their minds: Legal theory in neural networks. *Artificial Intelligence and Law*, 7(2):129–151, 1999.
- Mohammad Badiul Islam and Guido Governatori. Rulers: a rule-based architecture for decision support systems. *Artificial Intelligence and Law*, 26(4):315–344, 2018.
- Andrew Jones and Marek Sergot. Deontic logic in the representation of law: Towards a methodology. *Artificial Intelligence and Law*, 1(1):45–64, 1992.
- Grant Lamond. Do precedents create rules? *Legal Theory*, 11(1):1–26, 2005.
- Marc Lauritsen. On balance. *Artificial Intelligence and Law*, 23(1):23–42, 2015.
- Ronald Leenes. Burden of proof in dialogue games and Dutch civil procedure. In *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, pages 109–118, 2001.
- Fenrong Liu, Alessandra Marra, Paul Portner, and Frederik Van De Putte, editors. *Deontic Logic and Normative Systems, DEON 2020/2021*. College Publications, 2021.
- Arno R Lodder. *DiaLaw: on legal justification and dialogical models of argumentation*. Kluwer Academic Publishers, Dordrecht, 1999.
- Arno R Lodder, Thomas F. Gordon, The Pleadings Game — An Artificial Intelligence Model of Procedural Justice. *Artificial Intelligence and Law*, 8(2/3): 255–264, 2000.

- Ronald P Loui and Jeff Norman. Rationales and argument moves. *Artificial Intelligence and Law*, 3(3):159–189, 1995.
- Geraldine Mackenzie, Andrew Vincent, and John Zeleznikow. Negotiating about charges and pleas: balancing interests and justice. *Group Decision and Negotiation*, 24(4):577–594, 2015.
- James D Mackenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, pages 117–133, 1979.
- Juliano Maranhão. Refinement: A tool to deal with inconsistencies. In *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, pages 52–59, 2001.
- L Thorne McCarty. Some arguments about legal arguments. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pages 215–224, 1997.
- L.Thorne McCarty. Intelligent legal information systems: problems and prospects. In C. Campbell, editor, *Data processing and the Law*, pages 125–151. Sweet and Maxwell, London, 1984.
- Shah Jahan Miah, Jacqueline Blake, and Don Kerr. Meta-design knowledge for clinical decision support systems. *Australasian Journal of Information Systems*, 24:1–26, 2020.
- David Miller. Justice. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2021.
- Sanjay Modgil and Henry Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013.
- Sanjay Modgil and Henry Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument and Computation*, 5(1):31–62, 2014.
- Bernard Moulin, Hengameh Irandoust, Micheline Bélanger, and Gaëlle Desbordes. Explanation and argumentation capabilities: Towards the creation of more persuasive agents. *Artificial Intelligence Review*, 17(3):169–222, 2002.
- James O’ Neill, Paul Buitelaar, Cecile Robin, and Leona O’ Brien. Classifying sentential modality in legal language: a use case in financial regulations, acts and directives. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law*, pages 159–168, 2017.
- Truong-Son Nguyen, Le-Minh Nguyen, Satoshi Tojo, Ken Satoh, and Akira Shimazu. Recurrent neural network-based models for recognizing requisite and effectuation parts in legal texts. *Artificial Intelligence and Law*, 26(2):169–199, 2018.
- John L Pollock. How to reason defeasibly. *Artificial Intelligence*, 57(1):1–42, 1992.
- Henry Prakken. From logic to dialectics in legal argument. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 165–174, 1995.
- Henry Prakken. *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law*. Springer Science and Business Media B.V., Dordrecht, 1997a.
- Henry Prakken. Dialectical proof theory for defeasible argumentation with defeasible priorities (preliminary report). In *ModelAge Workshop on Formal Models of Agents*, pages 202–215. Springer, 1997b.
- Henry Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the 10th International Conference on Artificial Intelligence and Law*, pages 85–94, 2005.
- Henry Prakken. A formal model of adjudication dialogues. *Artificial Intelligence and Law*, 16(3):305–328, 2008.

- Henry Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.
- Henry Prakken. Reconstructing Popov v. Hayashi in a framework for argumentation with structured arguments and Dungean semantics. *Artificial Intelligence and Law*, 20(1):57–82, 2012.
- Henry Prakken. On the problem of making autonomous vehicles conform to traffic law. *Artificial Intelligence and Law*, 25(3):341–363, 2017.
- Henry Prakken. Modelling accrual of arguments in ASPIC+. In *Proceedings of the 17th International Conference on Artificial Intelligence and Law*, pages 103–112, 2019.
- Henry Prakken. A formal analysis of some factor-and precedent-based accounts of precedential constraint. *Artificial Intelligence and Law*, 29(4):559–585, 2021.
- Henry Prakken and Giovanni Sartor. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law*, 4(3-4), 1996.
- Henry Prakken and Giovanni Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6(2-4):231–287, 1998.
- Henry Prakken and Giovanni Sartor. Law and logic: A review from an argumentation perspective. *Artificial Intelligence*, 227:214–245, 2015.
- Henry Prakken, Adam Wyner, Trevor Bench-Capon, and Katie Atkinson. A formalization of argumentation schemes for legal case-based reasoning in ASPIC+. *Journal of Logic and Computation*, 25(5):1141–1166, 2015.
- Chris Reed and Tim Norman, editors. *Argumentation Machines. New Frontiers in Argument and Computation*. Kluwer Academic Publishers, Dordrecht, 2004.
- Adam Rigoni. An improved factor based approach to precedential constraint. *Artificial Intelligence and Law*, 23(2):133–160, 2015.
- Edwina L Rissland. Examples in legal reasoning: Legal hypotheticals. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 90–93, 1983.
- Edwina L Rissland. From UUM and CEG to CBR and ICAIL: A journey in AI and Law. In *From Knowledge Representation to Argumentation in AI, Law and Policy Making. A Festschrift in Honour of Trevor Bench-Capon on the Occasion of his 60th Birthday*, pages 191–212. College Publications, London, 2013.
- Edwina L Rissland and Kevin D Ashley. A case-based system for Trade Secrets law. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 60–66, 1987.
- Edwina L Rissland and David B Skalak. Cabaret: Statutory interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, 34:839–887, 1991.
- Edwina L Rissland, Eduardo M Valcarce, and Kevin D Ashley. Explaining and arguing with examples. In *Proceedings of the 4th AAAI Conference on Artificial Intelligence*, pages 288–294, 1984.
- Edwina L Rissland, David B Skalak, and M Timur Friedman. Bankxx: Supporting legal arguments through heuristic retrieval. *Artificial Intelligence and Law*, 4(1): 1–71, 1996.
- Edwina L Rissland, David B Skalak, and M Timur Friedman. Evaluating a legal argument program: The bankxx experiments. *Artificial Intelligence and Law*, 5(1):1–74, 1997.
- Giovanni Sartor. Normative conflicts in legal reasoning. *Artificial Intelligence and Law*, 1(2-3):209–235, 1992.

- Giovanni Sartor, Michał Araszkiewicz, Katie Atkinson, Trevor Bench-Capon, Floris Bex, Tom van Engers, Enrico Francesconi, Henry Prakken, and Giovanni Sileno. Thirty years of AI and Law: The second decade. *Artificial Intelligence and Law*, 30(4):This issue, 2022.
- Marek J Sergot, Fariba Sadri, Robert A Kowalski, Frank Kriwaczek, Peter Hammond, and H Terese Cory. The British Nationality Act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.
- Marek J Sergot, A S Kamble, and K K Bajaj. Indian central civil service pension rules: A case study in logic programming applied to regulations. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 118–127, 1991.
- Esther Shein. Converting laws to programs. *Communications of the ACM*, 65(1): 15–16, 2021.
- David M Sherman. A Prolog model of the Income Tax Act of Canada. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 127–136, 1987.
- Guillermo Simari and Ronald P. Loui. A mathematical treatment of defeasible reasoning and its applications. *Artificial Intelligence*, 53:125–157, 1992.
- Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2): 51–53, 2007.
- David B Skalak and Edwina L Rissland. Arguments and cases: An inevitable intertwining. *Artificial Intelligence and Law*, 1(1):3–44, 1992.
- Cor Steging, Silja Renooij, and Bart Verheij. Discovering the rationale of decisions: towards a method for aligning learning and reasoning. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pages 235–239, 2021.
- Andrew Stranieri and John Zeleznikow. Split Up: The use of an argument based knowledge representation to meet expectations of different users for discretionary decision making. In *Proceedings of AAI/IAAI 1998*, pages 1146–1151, 1998.
- Andrew Stranieri, John Zeleznikow, Mark Gawler, and Bryn Lewis. A hybrid rule–neural approach for the automation of legal reasoning in the discretionary domain of family law in Australia. *Artificial Intelligence and Law*, 7(2):153–183, 1999.
- Stephen E Toulmin. *The uses of argument*. Cambridge University Press, 1958.
- H. Jaap van den Herik. *Kunnen Computers Rechtspreken? (Can Computers Decide Legal Cases?)*. Gouda Quint, Arnhem, 1991.
- Bart Verheij. Deflog: on the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation*, 13(3):319–346, 2003a.
- Bart Verheij. Dialectical argumentation with argumentation schemes: An approach to legal logic. *Artificial intelligence and Law*, 11(2):167–195, 2003b.
- Bart Verheij. *Virtual Arguments. On the Design of Argument Assistants for Lawyers and Other Arguers*. T.M.C. Asser Press, The Hague, 2005.
- Bart Verheij. Formalizing arguments, rules and cases. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law*, pages 199–208, 2017.
- Bart Verheij. Artificial Intelligence as Law. *Artificial Intelligence and Law*, 28(2): 181–206, 2020.
- Bart Verheij, Jaap Hage, and H. Jaap van den Herik. An integrated view on rules and principles. *Artificial Intelligence and Law*, 6(1):3–26, 1998.

- Serena Villata, Michał Araszkiwicz, Kevin Ashley, Trevor Bench-Capon, L Karl Branting, Jack G Conrad, and Adam Wyner. Thirty years of AI and Law: The third decade. *Artificial Intelligence and Law*, 30(4):This issue, 2022.
- Pepijn Visser. *Knowledge specification for multiple legal tasks; a case study of the interaction problem in the legal domain*. Kluwer Law International, 1995.
- Pepijn Visser and Trevor Bench-Capon. A comparison of four ontologies for the design of legal knowledge systems. *Artificial Intelligence and Law*, 6(1):27–57, 1998.
- George Henrik von Wright. Deontic logic. *Mind*, 60:1–15, 1951.
- Gerard Vreeswijk. Abstract argumentation systems. *Artificial Intelligence*, 90:225–279, 1997.
- Gerard Vreeswijk and Henry Prakken. Credulous and sceptical argument games for preferred semantics. In *European Workshop on Logics in Artificial Intelligence*, pages 239–253. Springer, 2000.
- Douglas Walton. *Argumentation schemes for presumptive reasoning*. Lawrence Erlbaum Associates, 1996.
- Douglas Walton. Is there a burden of questioning? *Artificial Intelligence and Law*, 11(1):1–43, 2003.
- Douglas Walton. The basic slippery slope argument. *Informal Logic*, 35(3):273–311, 2015.
- Douglas Walton, Chris Reed, and Fabrizio Macagno. *Argumentation Schemes*. Cambridge University Press, Cambridge, 2008.
- Maya Wardeh, Trevor Bench-Capon, and Frans Coenen. Padua: a protocol for argumentation dialogue using association rules. *Artificial Intelligence and Law*, 17(3):183–215, 2009.
- Adam Wyner and Trevor Bench-Capon. Argument schemes for legal case-based reasoning. In *Proceedings of JURIX 2007*, pages 139–149. Citeseer, 2007.
- Zhengchuan Xu and Yufei Yuan. Principle-based dispute resolution for consumer protection. *Knowledge-Based Systems*, 22(1):18–27, 2009.
- John L Yearwood and Andrew Stranieri. The generic/actual argument model of practical reasoning. *Decision support systems*, 41(2):358–379, 2006.
- John Zeleznikow and Andrew 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*, pages 185–194, 1995.
- John Zeleznikow and Andrew Stranieri. Knowledge discovery in the Split-Up project. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pages 89–97, 1997.
- John Zeleznikow and Andrew Stranieri. Split-Up: an intelligent decision support system which provides advice upon property division following divorce. *International Journal of Law and Information Technology*, 6(2):190–213, 1998.
- John Zeleznikow, Andrew Stranieri, and Mark Gawler. Project report: Split-Up - a legal expert system which determines property division upon divorce. *Artificial Intelligence and Law*, 3:267, 1995.
- Heng Zheng, Davide Grossi, and Bart Verheij. Hardness of case-based decisions: a formal theory. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law*, pages 149–158, 2021a.
- Lucia Zheng, Neel Guha, Brandon R Anderson, Peter Henderson, and Daniel E Ho. When does pretraining help? Assessing self-supervised learning for law and the CaseHOLD dataset of 53,000+ legal holdings. In *Proceedings of the 18th*

International Conference on Artificial intelligence and Law, pages 159–168, 2021b.